NON-COMMUTATIVE ALGEBRA

MOSHE KAMENSKY

Abstract. Notes from a course on non-commutative algebra, BGU, Spring 2017

Contents

1. Introduction

1.1. **Representations.** Let G be a group. A fruitful approach to studying the properties of G is via its action on various sets. If X is a set endowed with a G-action (i.e., a G-*set*), and k is a field, the set k^X of all functions from X to k forms a k-algebra with pointwise addition and multiplication, and the action of G on X induces a G-action (by algebra automorphisms) on k^X, given by $g(f)(x) = f(g^{-1}x)$ (for $f \in k^X$ and $x \in X$).

G-set

Every element x of X determines a k-algebra map $e_x : k^X \to k$, the evaluation (projection) at x, and when X is finite, $x \mapsto e_x$ is a bijection from X to the set of k-algebra maps from k^X to k. More generally, maps $T : Y \to X$ between finite sets correspond bijectively to k-algebra maps $k^X \to x^Y$, so the properties of the action of G on X are completely reflected in its action on k^X and vice versa (for example, invariant subsets of X correspond to invariant ideals in k^X).

Thus, the transition from G-sets to algebras is not very useful: it is simply a change of language. The situation is different if we *forget* some of the structure, and regard k^X as a k-vector space. We thus obtain an example of a *representation* of G: a k-vector space endowed with a linear G-action. It seems that by doing so we lost information, but we gained in flexibility of the relations between object. For example, if $X = G$ with the regular action of G on itself, X clearly has no interesting invariant subsets, and therefore k^X has no interesting invariant ideals, but we will see that it usually has interesting invariant subspaces.

representation

One benefit of passing to linear representations is that the set of endomorphisms of a linear space V is itself a linear space. Thus, if G acts on V and $g, h \in G$, there is a natural way to associate an endomorphism of V to the formal sum $g + h$. Likewise, if $a \in k$, we obtain an endomorphism ag of V. Thus, we may associate a linear endomorphism of V to every element of the vector space $k[G]$ spanned by the elements of G. If we define a product on $k[G]$ via the group operation on the basis elements, we obtain the *group algebra* of G, and V becomes a module over it. Conversely, a module over $k[G]$ restricts to a representation of G, and maps of modules correspond to maps of representations. So we translated the study of representations of G to the study of modules over its group algebra $k[G]$. Thus, one of our motivating goals will be to *describe the category of modules over the group algebra* $k[G]$.

group algebra

1.2. **Semisimple algebras.** To explain the kind of description we are looking for, consider the case of the trivial group. Then, we are interested in the category of (finite dimensional) vector spaces. If V is such a space, and U

is a subspace, there is another subspace $W \subseteq V$ such that $W \oplus U \xrightarrow{\sim} V$ (we say that the category of vector spaces is *semi-simple*). By repeatedly taking proper subspaces, we may therefore write V as a finite direct sum of vector spaces that are *simple*, in the sense that they have no proper subspaces (finite dimension means that the sum is finite). Such a decomposition is unique, and the summands are all isomorphism to the field.

semi-simple

simple

For finite groups G, *Maschke's Theorem* implies that (under suitable assumptions on the field), a similar statement holds for modules over $k[G]$, except that the simple summands need not be isomorphic. An algebra with this property is called (Artinian-)*semisimple*. The structure of modules over such an algebra is obtained rather immediately from the structure of the algebra itself. Hence our goal is transformed into the task of *describing the structure of semisimple rings*.

The description is a classical theorem of Wedderburn, and will provide us with much information about the representations of G.

1.3. **Morita equivalence.** Since our motivating question concerns the category of modules over a ring, rather than the ring itself, it is natural to ask: could two essentially different algebras have the same category of modules? If the answer is positive, we could potentially benefit from switching between the two rings. To make the question precise, we will need to define precisely the meaning of the statement that "the categories are the same" or, more precisely, equivalent. Two such algebras are called *Morita equivalent*.

Morita equivalent

Many properties of an algebra can be defined in terms of its category of modules. Any such property is thus preserved across Morita equivalence. For commutative algebras, all properties are like that: Any two Morita equivalent commutative algebras are isomorphic. However, commutativity itself is not preserved under Morita equivalence: Our main result here will be that any ring R is Morita equivalent to the ring $M_n(R)$ of $n \times n$ matrices with entries in R ($n > 0$).

1.4. **Central simple algebras.** Wedderburn theory identifies the building blocks of semisimple k-algebras (k a field) to be matrix algebras $M_n(D)$, where D is a *division ring* (non-commutative field) over k. Hence to complete the classification, we need to study division algebras over k. However, the class of division algebras is sometimes inconvenient to work with directly: For example, the base change of a division algebra to a field extension of k need no longer be a division algebra. So we replace division algebras by the larger class of *central simple algebras*. On this class it is convenient to define directly certain operations, such as tensor products. Identifying a division algebra with the (Morita equivalent!) algebra of matrices over it, one obtains a group called the *Brauer group* of the field k. This is an invariant of the field, that plays an important role in number theory and geometry.

division ring

central simple algebras

Brauer group

1.5. **Basic definitions and notations.** By a *ring* we mean an associative

ring

ring with identity. This is mostly for didactic reasons, and to conform with what seems to be the standard agreement: many of the definitions and results hold for rings without identity, and we sometimes mention (and use) it. A *ring homomorphism* must be unital, i.e., map 1 to 1. If k is a ring, a *k-ring* (or a *ring over k*) is a ring R along with a prescribed ring homomorphism from k to R. A *k-algebra* (or an *algebra over k*) is a k-ring R where the image of k lies in the centre of R (the *centre* of R is the subring $Z(R)$ of R consisting of elements that commute with all other elements; in particular, it is commutative, and so is k). We will usually identify elements of k with their image under this map.

ring homomorphism
k-ring
ring over k
k-algebra
algebra over k
centre
$Z(R)$

If A is a ring, the *opposite ring* A° is the ring with the same underlying additive group as A, but with product reversed. Sometimes we wish to consider only the multiplicative monoid structure on A, which we denote by $A^\times$. A *left inverse* of an element $x \in A$ is an element $y \in A$ satisfying $yx = 1$. Such an element need not be unique, and need not be a right inverse. However, if x has both a left and a right inverse, then they coincide, and are unique (exercise!). An element that has a two sided inverse is called *invertible* or a *unit*. The group of all units of A is denoted by A^*.

opposite ring
A°
$A^\times$
left inverse
invertible

A *module* over a ring A is the analogue of a vector space, i.e., an abelian group M, along with an action map $a : A \times M \to M$ which is additive in each factor, and which a monoid action, i.e., $a(xy, m) = a(x, a(y, m))$ and $a(1, m) = m$ for all $x, y \in A$ and $m \in M$. More verbosely, this a *left module*, while a *right module* is a left module over A°. As usual, we omit a and write cm for $a(c, m)$ when possible. A map $T : M \to N$ of A-modules is a map of the underlying abelian groups that commutes with A: $T(cm) = cT(m)$ for all $c \in A$ and $m \in M$.

unit
A^*
module
left module
right module

If V is a module over a ring k, we denote by $\mathrm{End}_k(V)$ the ring of k-module maps from V to itself (the ring structure is given by addition of endomorphisms and composition). Note that this is a k-algebra when k is commutative.

1.6. **References.** A large portion of the course is based on the notes [Cla], and on the book [FD93] (useful, with many exercises, but beware of mistakes!). Other general references include the notes [Art], and the books [Row91] and [Her68] (which uses different terminology). Additional resources on specific topics are mentioned as they come.

2. Wedderburn–Artin theory

2.1. **Representations of finite groups.** Let G be a group. One approach to studying the group G is via its action on various objects. For example, Cayley's theorem asserts that any group is isomorphic to a subgroup of a symmetric group, via its *regular action* on itself. While it may seem that the simplest actions to consider are on sets, it is in fact beneficial to linearise such actions, and consider *linear* actions on vector-spaces. The definition actually makes sense more generally for a *monoid* (which is, by definition, a

monoid

set with an associative operation and a two-sided unit):

Definition 2.1.1. Let G be a monoid. A *representation* of G over a field k consists of a vector-space V over k, along with a map of monoids $r : G \to \mathrm{End}_k(V)$, i.e., $r(gh) = r(g) \circ r(h)$ for all $g, h \in G$, and $r(1)$ is the identity on V.

representation

A map from a representation (V_1, r_1) to another representation (V_2, r_2) of G (over the same field k) is a linear map $T : V_1 \to V_2$ such that $T \circ r_1(g) = r_2(g) \circ T$ for all $g \in G$.

As usual with actions, we shall abbreviate $r(g)(v)$ by gv when appropriate. Thus a representation is a left action on a vector space, where the map $v \mapsto gv$ is linear for all $g \in G$, and a map of representations is a map of actions which is linear.

Exercise 2.1.2. Let $T : U \to V$ be a map of representations. Show that the kernel and the image of T (as a linear map) are invariant under G. Show that if T is invertible (as a linear map), then the inverse is a map of representations.

From now on we fix a field k. A G-set (i.e., a set with a G-action) determines a representation over k as follows:

Example 2.1.3. For a set X, let k^X be the vector space of all functions from X to k. If $a : G \to \mathrm{Sym}(X)$ is a right action of G on X, then k^X becomes a representation of G via $r(g)(f)(x) = f(xg)$ (for $g \in G$, $f \in k^X$ and $x \in X$). □

Exercise 2.1.4. (1) Verify that this is so

(2) Show that if $u : X \to Y$ is a map of (right) G-sets, composition with u determines a map of representations $u^* : k^Y \to k^X$. Describe its kernel.

Assume that X is a G-invariant subset of a G-set Y. Applying the exercise to the inclusion, we obtain a sub-representation of k^Y (the kernel). Thus, invariant subsets determine invariant subspaces, but not all invariant subspaces are of this form:

Example 2.1.5. Assume that G is a non-trivial group, and let $X = G$ be the G-set with the (right) regular action of G on itself. Then X has no non-trivial invariant subsets, but the subspace of k^X consisting of constant functions is a 1-dimensional invariant (proper) subspace. □

Thus, the world of representations is, in some sense, richer: we may decompose the representation associated to a G-set, even if not the G-set itself.

Assume G is a group. If X is a G-set, and $Y \subseteq X$ is invariant, its complement Z is invariant as well. The complement of a linear subspace is never a subspace, but the subspaces corresponding to Y and Z form a direct-sum decomposition of k^X. We have seen above that not every invariant subspace of k^X is of this form, so it is reasonable to ask: Is every invariant subspace a direct summand? More generally, given a representation V and an invariant subspace U, is there an invariant subspace W such that $W \oplus U = V$? For finite groups, this is described by Maschke's theorem:

Theorem 2.1.6 (Maschke's Theorem)**.** *Let G be a group of (finite) order n, and let k be a field of characteristic p. The following are equivalent:*

(1) Any invariant subspace of any representation is a direct summand
(2) Any invariant subspace of k^G (the regular representation) is a direct summand
(3) p does not divide n

To give the proof, we first recall, from linear algebra, how to detect a subspace. If U is a linear subspace of a vector space V, direct sum decomposition for vector spaces implies that there is a linear map $T : V \to V$ whose image is U, and such that $T^2 = T$ (i.e., T is the identity when restricted to its image). Such a map is said to be an *idempotent*. Conversely, an idempotent T determines a direct sum decomposition consisting of the kernel and the image. By Exercise 2.1.2, this will be a direct sum of representations if T is a map of representations (i.e., commutes with the action of G).

idempotent

Proof of Theorem 2.1.6.

(2) $\implies$ (3)**:** Assume that p divides n, and let U be the subspace of constant functions in k^G. We identify an element of U with its constant value in k, and thus U with k. A complement to U thus corresponds to a map $T : k^G \to k$ which is invariant under the action of G and the identity on U. In particular, for all $v \in k^G$, we have $T(\sum_{g \in G} gv) = |G|T(v) = 0$ in k. Taking v to be the characteristic function of the identity in G, we have $\sum_{g \in G} gv = 1$, contradicting the assumption that $T(1) = 1$.

(3) $\implies$ (1)**:** Let U be an invariant subspace of V, and let T_0 be an idempotent with image U. Define $T : V \to V$ via

$$T(v) = \frac{1}{n} \sum_{g \in G} g^{-1} T_0(gv)$$

Then for all $v \in V$ and $h \in G$ we have

$$T(hv) = \frac{1}{n} \sum_{g \in G} g^{-1} T_0(ghv) = \frac{1}{n} \sum_{g' \in G} hg'^{-1} T_0(g'v) = hT(v)$$

where we substitute $g' = gh$. Also, for $v \in U$ we have $gv \in U$ and thus

$$T(v) = \frac{1}{n} \sum_{g \in G} g^{-1} T_0(gv) = \frac{1}{n} \sum_{g \in G} g^{-1} gv = v$$

So we found an idempotent that commutes with the action of G, and whose image is U.

(1) $\implies$ (2)**:** is trivial. □

2.1.7. *The group algebra.* One advantage of working with vector spaces (as opposed to sets) is that the set $\mathrm{Hom}_k(U, V)$ of linear maps between two spaces is again a vector space. Hence, if a representation assigns linear endomorphisms of V to elements $g, h \in G$, we have a natural linear endomorphism to assign to a formal linear combination $ag + bh$. This leads to the following definition:

Definition 2.1.8. Let G be a monoid, and let k be a field. The *group algebra*[1] $k[G]$ of G over k is the k-algebra whose underlying vector-space is the k-space spanned by the elements of G, and whose multiplication is extended by linearity from the product on G.

group algebra
$k[G]$

The inclusion of G in $k[G]^\times$ is a map of monoids into the multiplicative monoid of $k[G]$, which is universal in the following sense:

Exercise 2.1.9.
1. Let A be a k-algebra. Show that restriction determines a bijection between k-algebra maps from $k[G]$ to A and monoid maps from G to $A^\times$.
2. Conclude that the data of a representation on the vector-space V is equivalent that of a (left) $k[G]$ module structure on it (consider $A = \mathrm{End}_k(V)$ in the previous part).

End of lecture 1, Mar 13

2.2. **Categories and functors.** Much of the structure of a mathematical object is captured by the maps to or from that object that respect its structure. For example, the last exercise implies (with a bit more work, see below) that the group algebra of G can be characterised as the unique algebra H such that algebra maps from H correspond to monoid maps from G. Thus, one may hope to recover said structure from the knowledge of the maps that respect it. The most direct formal incarnation of this principle is encoded in the notion of a category:

Definition 2.2.1. A *category* $\mathcal{C}$ consists of the following data:

category

1. A collection of *objects* $\mathbf{Ob}_\mathcal{C}$ **objects**
2. For any two objects $X, Y \in \mathbf{Ob}_\mathcal{C}$ a set $\mathbf{Mor}_\mathcal{C}(X, Y)$ of *morphisms* from X to Y. **morphisms**
3. For any three objects X, Y, Z a *composition map* **composition map**

$$\circ : \mathbf{Mor}(Y, Z) \times \mathbf{Mor}(X, Y) \to \mathbf{Mor}(X, Z),\ \langle f, g\rangle \mapsto f \circ g$$

such that $f \circ (g \circ h) = (f \circ g) \circ h$ and for every object X the set $\mathbf{Mor}(X, X)$ includes an element 1_X with $1_X \circ f = f$ and $g \circ 1_X = g$, whenever the compositions make sense.

[1] Of course, the term "monoid algebra" would be more precise, but we will usually use the construction for groups, and "group algebra" is more standard.

We often write $f : X \to Y$ in place of $f \in \mathbf{Mor}_{\mathcal{C}}(X, Y)$.

Example 2.2.2. The category $\mathcal{S}et$ whose objects are all sets, and whose morphisms are usual maps between sets (with the usual composition). We could instead consider only injective (or only surjective) maps, or all relations (with composition of relations). □

Example 2.2.3. If R is a ring, the category $R - \mathcal{M}od$ of R-modules, with R-module maps between them. We will spend much time studying this example. Likewise, we have the categories $\mathcal{G}p$ of groups, $\mathcal{A}lg$ of rings, $\mathcal{T}op$ of topological spaces, etc. All with their usual morphisms. □

pointed sets

Example 2.2.4. The category $\mathcal{S}et_*$ of *pointed sets* has objects of the form (X, x), where X is a set and $x \in X$. A morphism from (X, x) to (Y, y) is a map of sets $f : X \to Y$ such that $f(x) = y$. The category $\mathcal{T}op_*$ of pointed topological spaces is defined similarly. □

Example 2.2.5. Let X be a topological space. We may associate to X a category $\pi_1(X)$ as follows: the objects are the points of X, morphisms from $x \in X$ to $y \in X$ are homotopy classes of paths from x to y, and composition is (induced by) concatenation of paths. Thus, the endomorphisms of each object x is the fundamental group based at x. □

opposite category
$\mathcal{C}^\circ$

Example 2.2.6. If $\mathcal{C}$ is a category, the *opposite category* $\mathcal{C}^\circ$ is the category with the same objects as $\mathcal{C}$, but with morphisms inverted: $\mathbf{Mor}_{\mathcal{C}^\circ}(X, Y) = \mathbf{Mor}_{\mathcal{C}}(Y, X)$ (and composition with switched arguments). It follows that every notion in category theory has a *dual notion*, which is the original notion applied to the opposite category. □

Example 2.2.7. If X is an object of $\mathcal{C}$, composition gives $\mathbf{Mor}(X, X)$ the structure of a monoid. This process identifies monoids with categories with one object. The opposite category corresponds, in this case, to the opposite monoid. □

Example 2.2.8. On the other extreme, if R is a binary relation on a set A, we may try to construct a category whose objects are the elements of A, and with $\mathbf{Mor}(a, b)$ consisting of one element if $\langle a, b\rangle \in R$, and empty otherwise. This construction yields a category precisely if R is a pre-order (i.e., reflexive and transitive), and conversely, a (small) category where all $\mathbf{Mor}$ sets have size at most 1 determines a pre-order on the set of objects. In particular, every set can be viewed as a category, where the only morphisms are identities. □

Some properties of maps can be expressed through composition alone, and thus make sense for morphisms in any category. For example, a morphism $f : X \to Y$ is an *isomorphism* if it has a two-sided inverse, i.e., a morphism $g : Y \to X$ with $f \circ g = 1_Y$ and $g \circ f = 1_X$.

isomorphism

We mentioned above that familiar structures are encoded by morphisms. We illustrate it with a few simple examples. The empty set can be characterised as an object of $\mathcal{S}et$ that has precisely one morphism to each object.

An object in a category with this property is called an *initial object*. Thus, we automatically obtain the analogue of the empty set in each category (which need not exist):

initial object

Exercise 2.2.9.
1. Show that if X and Y are two initial objects then there is a unique morphism from X to Y, which is an isomorphism. Thus, the initial object (if exists) is essentially unique.
2. Compute the initial object in each of the examples above (or show that it does not exist).
3. The dual notion to the initial object is the *final object*. Compute the final object in $\mathcal{S}et$ and in the other examples.

final object

As another example, consider two sets X and Y. If we view these sets as embedded in a bigger set (or "universe"), it makes sense to consider their union, but if no such embedding is given this is meaningless, since we cannot determine when an element of X is equal to an element of Y. However, it does make sense to take their *disjoint* union. This is a set $X\coprod Y$ with the following property: defining a map from $X\coprod Y$ to some other set Z is canonically equivalent to defining its restrictions to X and to Y. This is phrased almost in purely categorical terms, but we have to understand what restrictions are. For that, note that we are given maps $i_X : X \to X\coprod Y$ and $i_Y : Y \to X\coprod Y$, and the restriction of a function f from $X\coprod Y$ to X is nothing but the composition $f \circ i_X$.

Thus, given any set Z with maps $f : X \to Z$ and $g : Y \to Z$, there is a unique way to amalgamate them to a map $h : X\coprod Y \to Z$ so that $h\circ i_X = f$ and $h \circ i_Y = g$. This property completely determines the disjoint union, in a sense made precise below, and since it is given in terms of maps and their compositions, provides an analogue of the disjoint union in an arbitrary category:

Definition 2.2.10. A *coproduct* $X\coprod Y$ of two objects X, Y in a category $\mathcal{C}$ is an object W, along with morphisms $i_X : X \to W$ and $i_Y : Y \to W$, which is *universal*, i.e., for any object Z and morphisms $f : X \to Z$ and $g : Y \to Z$ there is a unique morphism $h : W \to Z$ with $h \circ i_X = f$ and $h \circ i_Y = g$.

coproduct
$X\coprod Y$

The dual notion is called a *product*, denoted $X \times Y$.

product
$X \times Y$

We will normally say that W is a coproduct of X and Y, even though the morphisms i_X, i_Y are part of the data. The coproduct and product of an arbitrary set of objects is defined analogously.

Exercise 2.2.11.
1. Show that if W and W' are both coproducts of two objects X and Y, there is a unique morphism from W to W' respecting the coproduct structure, and this morphism is an isomorphism. Hence, we may speak of *the* coproduct
2. Show that if the coproduct of X and Y exists, then so does the coproduct of Y and X, and they are uniquely isomorphic. Likewise, $(X\coprod Y)\coprod Z = X\coprod(Y\coprod Z)$ (i.e., if one side exists so does the other, and they are *uniquely* isomorphic).

(3) Compute the products and coproducts in the categories above (some of these are difficult!)

Exercise 2.2.12. If X, Y are objects of a category $\mathcal{C}$, let $\mathcal{C}_{X,Y}$ be the category whose objects are triples $\langle Z, f, g\rangle$ where Z is an object of $\mathcal{C}$, and $f : X \to Z$, $g : Y \to Z$ are morphisms in $\mathcal{C}$. A morphism from $\langle Z_1, f_1, g_1\rangle$ to $\langle Z_2, f_2, g_2\rangle$ is a morphism $h : Z_1 \to Z_2$ in $\mathcal{C}$, such that $h \circ f_1 = f_2$ and $h \circ g_1 = g_2$. Show that $\langle Z, f, g\rangle$ is an initial object of $\mathcal{C}_{X,Y}$ if and only if it is a coproduct in $\mathcal{C}$.

The philosophy of category theory entails that whenever we define a class of objects, we should also define what are the morphisms between them. We have defined categories, so what are the morphisms between them?

functor

Definition 2.2.13. Let $\mathcal{C}$ and $\mathcal{D}$ be categories. A *functor* from $\mathcal{C}$ to $\mathcal{D}$ consists of a map $\mathbf{F} : \mathbf{Ob}_{\mathcal{C}} \to \mathbf{Ob}_{\mathcal{D}}$ and for any two objects X, Y of $\mathcal{C}$ a map $\mathbf{F}_{X,Y} : \mathbf{Mor}_{\mathcal{C}}(X, Y) \to \mathbf{Mor}_{\mathcal{D}}(\mathbf{F}(X), \mathbf{F}(Y))$ such that $\mathbf{F}(g \circ h) = \mathbf{F}(g) \circ \mathbf{F}(h)$ and $\mathbf{F}(1_X) = 1_{\mathbf{F}(X)}$ for all X and all composable g, h.

contravariant functor

A *contravariant functor* from $\mathcal{C}$ to $\mathcal{D}$ is a functor from $\mathcal{C}$ to $\mathcal{D}^\circ$

Intuitively, a functor is an assignment of objects of type $\mathcal{D}$ to objects of type $\mathcal{C}$, in a canonical manner.

Example 2.2.14. If $\mathcal{G}p$ is the category of groups (and group homomorphisms), there is a functor from it that assigns to each group its underlying set, and to each homomorphism itself. This functor is often called the "forgetful" functor, since it forgets the group structure. There are, similarly, forgetful functors from rings to abelian groups, from pointed sets to sets, etc. □

Example 2.2.15. In the other direction, the process of constructing, from a set X, the free group $F(X)$ generated by X is a functor: a map of sets from X to Y determines a map from X to $F(Y)$ (since $Y \subset F(Y)$), and this map extends uniquely to a map of groups $F(X) \to F(Y)$. Other free constructions (free abelian group, vector space spanned by a set, free algebra, etc.) are also functorial, for similar reasons. □

The "free" constructions are clearly not inverse to the corresponding "forgetful" ones, but are nonetheless related, in a manner that will be described later.

Example 2.2.16. Every abelian group is, in particular, a group, and a map of abelian groups is the same as a map of groups. Therefore, we have an "inclusion" of the category of abelian groups in groups. In the other direction, for any group G there is an abelian group G^{ab}, along with a group homomorphism $G \to G^{ab}$, which is universal: any group homomorphism $G \to A$ into an abelian group A factors uniquely via a map $G^{ab} \to A$. G^{ab} is called the *abelianisation* of G, and can be consructed as G/G' where G' is the commutator subgroup of G. □

abelianisation

Exercise 2.2.17. Define the abelianisation on the level of morphisms, and show that this is indeed a functor (note that you *do not* need to use the construction of G^{ab})

Example 2.2.18. The fundamental group is a functor $\pi_1 : \mathcal{T}op_* \to \mathcal{G}p$ from pointed topological spaces to groups: a continuous map $f : X \to Y$ determines a group homomorphism $\pi_1(f) : \pi_1(X, x) \to \pi_1(Y, f(x))$, induced by composition[2]. Other homotopical invariants (e.g. Homology) are likewise functorial. □

Example 2.2.19. If G and H are monoids, viewed as categories with one object, a functor from G to H is the same as a map of monoids. If $\mathcal{C}$ is any category, a functor from G to $\mathcal{C}$ is (equivalent to) an object X of $\mathcal{C}$ along with an action of G on X. For example, if $\mathcal{C}$ is $\mathcal{S}et$, such a functor is a G-set, while if $\mathcal{C}$ is the category of vector spaces over a field k, it is a representation. □

Example 2.2.20. If k is a field, assigning to each vector space over k its dual is a contravariant functor from the category $\mathcal{V}ec_k$ of vector spaces over k to itself. □

Example 2.2.21. For any set X, let $\mathcal{P}(X)$ be its power set. If $f : X \to Y$ is a function, let $\mathcal{P}(f) : \mathcal{P}(Y) \to \mathcal{P}(X)$ be the map $Z \mapsto f^{-1}(Z)$. This is a contravariant functor from $\mathcal{S}et$ to itself. □

Example 2.2.22. For any field k, Example 2.1.3 attaches a vector space k^X to each set X, and Exercise 2.1.4 attaches a linear map u^* to each map of sets. This determines a contravariant functor from $\mathcal{S}et$ to $\mathcal{V}ec_k$ (and from G-sets to G-representations, for each monoid G). □

Exercise 2.2.23. Let $\mathbf{F} : \mathcal{C} \to \mathcal{D}$ be a functor. Show that if both $\mathcal{C}$ and $\mathcal{D}$ have an initial object (both denoted by 0), there is a canonical map $0 \to \mathbf{F}(0)$. Likewise, for any two objects X, Y there is a canonical map from $\mathbf{F}(X) \coprod \mathbf{F}(Y)$ to $\mathbf{F}(X \coprod Y)$, assuming these coproducts exist.

If the maps in the last exercise are isomorphisms, we say that $\mathbf{F}$ *preserves finite coproducts*. Note that if we apply this definition to contravariant functors, it means that $\mathbf{F}$ takes coproducts to the corresponding products, in terms of the original categories.

preserves finite coproducts

Exercise 2.2.24. Show that the functor from Exercise 2.2.22 preserves coproducts.

End of lecture 2, Mar 14

2.3. **Semisimple modules.** The property of representations in Mashcke's Theorem is captured by the following general definition:

Definition 2.3.1. Let R be a ring. An R-module M is a *semisimple module* if every submodule is a direct summand (i.e., if $N \subseteq M$ is a submodule, then there is a submodule $L \subseteq M$ such that $N \oplus L = M$)

semisimple module

A ring R is a *semisimple ring* if it is semisimple as a module over itself.

semisimple ring

[2] As explained above, we may avoid fixing a base point by attaching, to a space X, the category $\pi_1(X)$ of paths. This, once again, is a functor, i.e., each continuous map induces a functor between the corresponding categories. We shall not pursue this direction in this course

Example 2.3.2. Maschke's Theorem asserts that if G is a finite group then every $\mathbb{C}[G]$-module is semisimple. □

We will consider more examples after the following result.

Proposition 2.3.3. *Let R be a ring*

1. *If M is a semisimple R-module, then so is every submodule and every quotient of M*
2. *If $\{M_\alpha\}$ is a family of semisimple R-modules, so is their direct sum*

Before giving the proof, we note that the quotient module M/N has a categorical description, namely, it is a module L along with a map $t : M \to L$ such that $t(N) = 0$, and is universal for this property. As such, it is uniquely determined, up to a unique isomorphism. If $M = N \oplus K$ for some submodule K, the projection from M to K is easily seen to have this universal property, hence K is isomorphic to M/N (and more precisely, $\langle M/N, t\rangle$ is uniquely isomorphic to $\langle K, \pi\rangle$, where π is the projection). In particular, any two direct summands of N inside M are isomorphic.

Proof of Prop. 2.3.3.

(1) By the remarks above, we only need to prove for submodules. Let $N \subseteq M$ be a submodule, and let $K \subseteq N$. Since M is semisimple, N has a direct summand L in M, and $L+K$ has a direct summand E. We claim that the projection $\pi : M \to N$ restricts to an isomorphism from $K + E$ to N. Indeed, since $M = E \oplus L \oplus K$, the kernel L of π intersects $E \oplus K$ trivially, so π is injective when restricted to $E \oplus K$. On the other hand, if $n \in N$, writing $n = e + l + k$ according to the direct sum we have $n = \pi(n) = \pi(e + k) + \pi(l) = \pi(e + k)$, hence n is in the image.

(2) We first show the statement for two modules, in a stronger form: If $E \subseteq M \oplus N$, K is a direct summand of $E \cap M$ in M and L is a direct summand of $\pi(E)$, where $\pi : M \oplus N \to N$ is the projection, then $K + L$ is a direct summand of E.

Assume that $x \in (K + L) \cap E$. Then $\pi(x) \in \pi(E)$, but also $\pi(x) \in L$, since $K \subseteq M = \ker(\pi)$, so $\pi(x) = 0$. But then $x \in M$, hence $x \in K$ and also $x \in E \cap M$, so $x = 0$. To show that $M \oplus N = K + L + E$, it suffices to show that $M \subseteq K + L + E$ and $N \subseteq M + L + E$. The first holds since $M \subseteq K + E$. For the second, let $n \in N$. Then for some $e \in E$ we have $n = l + \pi(e) - e + e$, and $\pi(e) - e \in M$. This completes the case of two modules.

Now let $\{M_\alpha\}$ be an arbitrary collection of semisimple modules. We may assume them to be well-ordered, by an ordinal β. Let $E \subseteq \bigoplus_{\alpha<\beta} M_\alpha$, and set $M^\gamma = \bigoplus_{\alpha<\gamma} M_\alpha$, $E^\gamma = E \cap M^\gamma$. We will construct, by induction, direct summands F^γ of E^γ in M^γ, with the property that $F^\gamma \subseteq F^\delta$ for $\gamma < \delta$. Taking $\gamma = \beta$, this will prove the claim.

For successor stages, we apply the case of two modules, with M^γ and M_γ. When γ is limit, we set $F^\gamma = \bigcup_{\alpha<\gamma} F^\alpha$. Since $M^\gamma = \bigcup_{\alpha<\gamma} M^\alpha$, we have $F^\gamma \subseteq M^\gamma$.

Assume that $x \in F^\gamma \cap E^\gamma$. Then for some $\alpha < \gamma$ we have $x \in F^\alpha \cap E^\alpha$, hence $x = 0$ by induction. Likewise, if $m \in M^\gamma$, then $m \in M^\alpha$ for some $\alpha < \gamma$, so $m = e + f$ with $e \in E^\alpha \subseteq E^\gamma$ and $f \in F^\alpha \subseteq F^\gamma$. □

The converse of the first part need not be true: If M is a module with a semisimple submodule N such that the quotient M/N is semisimple, M need not be semisimple.

We now consider two examples from the commutative world.

Example 2.3.4. Let $R = \mathbb{Z}$. Then modules are abelian groups. A finitely generated abelian group a direct sum of cyclic groups whose order is a prime power, or infinite. According to the proposition, we may consider each summand separately. If A is such a summand, and $B \subseteq A$ is a non-trivial subgroup, it cannot be a direct summand by the uniqueness part of the classification. Hence, a finitely generated abelian group is semisimple if and only if it is a direct sum of groups of prime order. □

Example 2.3.5. Let $R = k[x]$, where k is an infinite field. We view the elements of R as functions on k. We recall that every ideal of R is generated by a single polynomial. If I and J are ideals generated by p and q, then pq belongs to their intersection. Hence, R itself is not semisimple.

If V is a vector space over k, the space M_V of functions from k to V is a module over R, which acts by multiplication of functions. Let $f \in M_V$ be a function. If f is non-zero at infinitely many points of k, then pf is not the zero function for any non-zero polynomial p. Hence, the submodule generated by f is isomorphic to R (by the isomorphism taking $1 \in R$ to f), so is not semisimple.

Thus, if $M \subseteq M_V$ is semisimple, then every element of M has *finite support*, i.e., is identically 0 outside a finite set. On the other hand, the submodule $M_0 \subseteq M_V$ consisting of all finitely supported functions is semisimple: it is the direct sum $\bigoplus_{a\in k} V_a$, where V_a is the module of functions supported only at a. Each V_a is isomorphic to V as a vector space over k (via the map $R \to k$ given by evaluation at a), and is thus semisimple.

We conclude that a submodule of M_V is semisimple if and only if it is contained in M_0. □

The preceding examples show that in the commutative world, semisimple modules correspond to the simplest geometric situation, that of a finite (discrete) set, and are thus not very interesting. On the other hand, for general rings the theory is much richer, and includes, for example, the representation theory (in characteristic 0) of all finite groups.

End of lecture 3, Mar 22

Corollary 2.3.6. *If a ring R is semisimple, then all its modules are semisimple*

The last proposition already shows that semisimple modules behave similarly to vector spaces. A vector space further has the property that it is a direct sum of very simple pieces: one-dimensional spaces. We have no notion of dimension, and no hope for such a simple description, but we have replacement for one-dimensional spaces:

simple module

Definition 2.3.7. A module M over R is a *simple module* if it has precisely two submodules (namely, 0 and M)[3]

Proposition 2.3.8. *A module is simple if and only if it is isomorphic to R/I, where I is a maximal (left) ideal*

While it is easy enough to give the definition and construct them, it is not clear how often simple modules are encountered. In fact:

Example 2.3.9. $\mathbb{Z}$ has no simple submodules (as a module over itself) □

However, the situation is different for semisimple modules:

Lemma 2.3.10. *Any non-zero semisimple module contains a simple submodule*

Proof. Let $a \in M$ be a non-zero element of a semisimple module M. By Zorn's Lemma, M has a submodule N which is maximal among submodule not containing a. By semisimplicity, it has a direct summand L. We claim that L is simple: otherwise, it has a non-trivial submodule K, and a direct summand (within L), E. At least one of K, L does not contain a, and therefore can be added to N, to contradict maximality. □

Proposition 2.3.11. *Let M be an R-module. The following are equivalent:*

(1) M is semisimple
(2) M is a direct sum of simple modules
(3) M is the sum of its simple submodules

Proof. (2) $\implies$ (1)**:** follows from Prop 2.3.3, since any simple module is semisimple.

(3) $\implies$ (1)**:** Let M' be the direct sum of all simple submodules of M. By the previous implication, M' is semisimple, and the sum of all simple submodules of M is a homomorphic image, so by Prop 2.3.3 is semisimple. By assumption, this sum is M.

(2) $\implies$ (3)**:** is trivial.

(1) $\implies$ (2)**:** Call a set X of simple submodules of M *independent* if their sum is direct. Let P be the set of all independent sets of simple submodules of M, ordered by inclusion. The union of a chain in P is again in P: If the sum of a union is not direct, there is a finite non-trivial equation that holds there, and this equation holds at a finite stage.

[3] Thus, 0 is semisimple but not simple

Applying Zorn's Lemma, let X be a maximal element of P, and $N = \bigoplus X \subseteq M$. By (1), N has a direct summand E. If $E \neq 0$, it has a simple submodule S, which is independent of X, contradicting the maximality of X. Thus $N = M$. □

Exercise 2.3.12. Show that any module M has a biggest semisimple submodule. This submodule is called the *socle* of M, denoted $\mathrm{soc}(M)$. Show that if $f : M \to N$ is a map of modules, then f restricts to a map from $\mathrm{soc}(M)$ to $\mathrm{soc}(N)$. Thus, soc is a functor from R-modules to semisimple R-modules. **socle**

If R is a field, there is just one isomorphism class of simple modules, namely, the class of R viewed as a module over itself. Since R is then simple, it is semisimple, and we get the familiar fact that every vector space is a direct sum of copies of R, i.e., every vector space has a basis. In the case of vector spaces, we also know that the number of copies needed is independent of the decomposition, and that this number determines the isomorphism class (i.e., there is a well defined notion of dimension, and two modules of the same dimension are isomorphic).

The same cannot literally be true in the general case, since we have non-isomorphic simple modules. However, we will see that this is the only obstacle: For a semisimple module M, the number of times a simple module S occurs in a decomposition is well defined, and the collection of these numbers, for all possible simple modules S, determines the isomorphism class of M. To show this, we need to discuss morphisms between semisimple modules. We start with the simple ones.

Proposition 2.3.13 (Schur's Lemma). *If $f : M_1 \to M_2$ is a map between simple modules, then f is either 0 or an isomorphism*

Exercise 2.3.14. Prove Schur's Lemma

To pass to general semisimple modules, we note that for any collection M_α of modules, the direct sum $\oplus_\alpha M_\alpha$ is naturally a submodule of the direct product $\prod_\alpha M_\alpha$ (the direct sum is identified with those elements of the product that have *finite support*, i.e., that are 0 except at a finite set of indices). Since the direct sum and the direct product are the categorical coproduct and product, respectively, in the category of modules, we may make use of the following observation:

Exercise 2.3.15. Show that given objects X_α and Y_β in a category $\mathcal{C}$, there is a natural bijection

$$\mathbf{Mor}_{\mathcal{C}}(\coprod_\alpha X_\alpha, \prod_\beta Y_\beta) = \prod_{\alpha,\beta} \mathbf{Mor}_{\mathcal{C}}(X_\alpha, Y_\beta)$$

(assuming, of course, that these coproduct and product exist; here α and β range over some fixed set)

Combining the last exercise with the observations before it, and Schur's Lemma, we obtain:

Corollary 2.3.16. *For any two collections S_α and T_β of simple modules, for $M = \oplus_\alpha S_\alpha$ and $N = \oplus_\beta T_\beta$ we have*

$$\operatorname{Hom}(M, N) \subseteq \prod_{X \text{ semisimple}} \operatorname{Hom}(\oplus_{s_X} X, \oplus_{t_X} X)$$

where s_X is the cardinality of the set of indices α for which S_α is isomorphic to X (and similarly for t_X). Furthermore, if $f \in \operatorname{Hom}(M, N)$ is injective, or an isomorphism, then so is each of its components on the right.

The last corollary reduces the question of isomorphism of two direct sums to the case of a single simple module. To complete the classification, we show that, just like in vector spaces, the isomorphism type is determined by the number of summands:

Proposition 2.3.17. *If S is a simple module, and $\oplus_{\alpha \in I} S$ is isomorphic to $\oplus_{\alpha \in J} S$, then I and J have the same cardinality.*

The proof is virtually identical to the special case of vector spaces over a field.

Proof. We first claim that if $T \subset \oplus_n S$ is a proper submodule (where n is finite), then T is isomorphic to $\oplus_m S$, where $m < n$. Consider the projection π on the first $n-1$ coordinates. Then T is isomorphic to $\pi(T) + \ker(\pi)$. Since $\ker(\pi) \subseteq S$ and S is simple, we have $\ker(\pi) = 0$ or $\ker(\pi) = S$. In the first case, T is isomorphic to the submodule $\pi(T)$ of a smaller sum, so we are done by induction. If $\pi(T)$ is a proper submodule of $\oplus_{n-1} S$, we are again done by induction. Hence we may assume that $\ker(\pi) = S$ and $\pi(T) = \oplus_{n-1} S$, but then T is not a proper submodule.

The claim proves the proposition in the case that one of the index sets is finite. Assume now that that $t : \oplus_I S \to \oplus_J S$ is an isomorphism, where I and J are infinite. Let $s \in S$ be a non-zero element, and let $e_\alpha \in \oplus_I S$, for $\alpha \in I$, be the element value s at α and 0 elsewhere. Since s generates S, the e_α generate $\oplus_I S$. Each $t(e_\alpha)$ has a finite support $J_\alpha \subseteq J$, so each element in the image of t has support in $J' = \cup_{\alpha \in I} J_\alpha$. On the other hand, since t is an isomorphism, $J' = J$. Since J' is a union of $|I|$ finite sets, we have $|J| \leq |I|$. Applying this in the other direction we get equality. □

To conclude, the isomorphism type of a semisimple module is determined by its simple summands. To state this more precisely, we introduce the following definition:

Definition 2.3.18. If M is a semisimple module, and S is a simple module, the *multiplicity* of S in M, denoted $m_S(M)$, is the number of copies of S that appear in the direct sum decomposition of M into simple modules (this is a cardinal).

multiplicity
$m_S(M)$

The preceding results state that the numbers $m_S(M)$ are well defined for each S and M (i.e., do not depend on the decomposition), as well as:

Corollary 2.3.19. *Semisimple modules M and N over a ring R are isomorphic if and only if $m_{R/I}(M) = m_{R/I}(N)$ for every maximal (left) ideal I.*

In the case when R is a field, 0 is the only (hence maximal) left ideal, and $m_R(M)$ is just the dimension of M as a vector space. Thus, the corollary is a direct generalisation of linear algebra.

Example 2.3.20. Continuing Example 2.3.5, each $a \in k$ corresponds to the maximal ideal generated by $x-a$, and the corresponding multiplicity $m_a(M)$ (for $M \subseteq M_0$) is the dimension of the space of functions in M that are supported on a. □

Exercise 2.3.21. Compute a direct sum decomposition of $k[S_3]$ as a module over itself (where $\mathrm{char}(k) \neq 2, 3$)

End of lecture 4, Mar 23

End of lecture 5, Mar 29

2.4. **Semisimple rings.** We now consider the case that R itself is semisimple. Recall that this is satisfied by the group algebra of a finite group over almost all characteristics, and that in this situation, any module is semisimple. Our goal is to understand the structure of R as a ring. We begin with a few observations about R as a module over itself, which hold for a general ring.

Proposition 2.4.1. *If $R = \cup_{\alpha \in A} M_\alpha$ where $M_\alpha \subseteq R$ is a submodule for each α, then $R = M_\alpha$ for some α*

Proof. For some α, $1 \in M_\alpha$. □

The module R has and additional special feature: it is also a *right* R-module. We now study the action of R on itself on the right. We first make a simple observation:

Proposition 2.4.2. *Let R be any ring. The action of R on itself on the right identifies R with $\mathrm{End}_R(R)^\circ$, the opposite of the endomorphism ring of R as a module over itself.*

Exercise 2.4.3. Prove the proposition

Corollary 2.4.4. *A left ideal in R is a two-sided ideal if and only if it is preserved by all endomorphisms of R (as a module)*

Proof. A left ideal is a two-sided ideal if and only if it is invariant under multiplication on the right by all elements, which coincides with the action of endomorphisms □

Here is the application we will need:

Corollary 2.4.5. *Let S be a simple module over a ring R, and let I_S be the sum of all ideals in R isomorphic to S. Then I_S is a two-sided ideal.*

Proof. I_S is a sum of left submodules, so is a left submodule. Any endomorphism of R will take an ideal isomorphic to S to another such (or to 0), so the sum is invariant under all endomorphisms. □

If R is semisimple, it is a sum of itself simple submodules, so we find that R is a direct sum of two-sided ideals (each the sum of a single isomorphism class of simple modules). This allows us to use the following result, known as the *Chinese remainder theorem*:

Proposition 2.4.6 (The Chinese Remainder Theorem)**.** *Let R be any ring, $I_1, \dots, I_n$ two-sided ideals such that $I_k + I_l = R$ for all $k \neq l$, and $\cap_j I_j = 0$. Let $R_k = R/I_k$, and let $\pi_k : R \to R/I_k$ be the projection. Then the combined map $\pi_1 \times \dots \times \pi_n : R \to \prod_i R_i$ is an isomorphism*

Proof. The map is clearly injective. To prove surjectivity, for each $k < n$ we may write $1 = e_k + f_k$, with $e_k \in I_k$ and $f_k \in I_n$. Taking the product over all $k < n$, we get

$$1 = \prod_k (e_k + f_k) = \prod_k e_k + r$$

where r is the sum of products, each of which consisting of at least one f_k. Writing $J = \cap_{k<n} I_k$, we thus have $r \in I_n$ and $\prod_k e_k \in J$, so $I_n + J = R$ and $I_n \cap J = 0$, and we are reduced to the case $n = 2$. The elements $(1,0), (0,1) \in R/I_1 \times R/I_2$ generate it as an R-module, and are the images of f_1, e_1, respectively, so we are done. □

Corollary 2.4.7. *Any semisimple ring is a finite product of rings having precisely one simple module (up to isomorphism)*

The product decomposition results in a decomposition of the whole theory:

Exercise 2.4.8. Let R and S be two rings. Let $\mathcal{C}$ be the category whose objects are pairs (M, N) where M is an R-module and N is an S-module, and where morphisms are pairs of module homomorphisms. Show that there is a functor $\mathbf{F}$ from the category $R \times S - \mathcal{M}od$ of module over $R \times S$ to $\mathcal{C}$, and a functor $\mathbf{G} : \mathcal{C} \to R \times S - \mathcal{M}od$, and isomorphisms $\alpha_M : M \to \mathbf{G}(\mathbf{F}(M))$ and $\beta_{(N,L)} : (N, L) \to \mathbf{F}(\mathbf{G}((N, L))$ for all M, N, L in the appropriate categories.

Remark 2.4.9. In terminology that will be defined in §5.1.2, this almost shows that the category of $R \times S$-modules is equivalent to $\mathcal{C}$. □

Because of the last exercise, we may concentrate on each factor separately. The factors can be characterized via the following definition:

faithful module
primitive ring

Definition 2.4.10. A module M over a ring R is a *faithful module* if for all $r \in R$, if $rM = 0$, then $r = 0$. A ring R is a *primitive ring* if it has a faithful simple module.

Exercise 2.4.11. Show that if $R = R_1 \times R_2$ is a non-trivial product of rings, then it cannot be primitive.

The notion of a primitive ring is interesting outside the context of semisimple rings. In the semisimple case, we have:

Proposition 2.4.12. *A semisimple ring R is primitive if and only if it has only one simple module (up to isomorphism)*

Proof. If R has only one simple module S, it is isomorphic (as a module) to S^n for some n. Hence if $rS = 0$ then $0 = rS^n = rR$, so $r = 0$ since $r = r1 \in rR$.

If R has more than one simple module, we had seen that it is a product of rings, so cannot be primitive by the last exercise. □

So we now concentrate on the case of primitive semisimple rings. We first consider an example.

Example 2.4.13. Let V be a vector space over a field k, and let $R = \operatorname{End}_k(V)$. Then R acts on V, and since the action is transitive on (non-zero) vectors, this is a simple module. Thus R is primitive.

If B is a basis for V, any map (of sets) $T_0 : B \to V$ extends to a unique linear endomorphism of V. Thus, as R-modules, R is isomorphic to V^B. In particular, R is semisimple if and only if B is finite, and in this case, V is the only simple module. □

We now wish to ask to which extent the last example is typical: is any primitive semisimple ring R necessarily like in the last example? We will see that the answer is yes, provided we allow the field to be non-commutative:

Definition 2.4.14. A *division ring* is a ring in which every non-zero element is invertible. **division ring**

A division ring is clearly primitive semisimple, but if it is not commutative, it is not of the form $\operatorname{End}_k(V)$ for a commutative field k. On the other hand, essentially all of linear algebra generalises to vector spaces (modules) over division rings, so rings of the form $\operatorname{End}_k(V)$ where k is a division ring are rather familiar as well.

We will later see that non-commutative division rings do exist. In the meantime, we note the following general method of producing division rings:

Exercise 2.4.15. Show that if S is a simple module over some ring R, then $\operatorname{End}_R(S)$ is a division ring

We may now prove the classification result:

Theorem 2.4.16. *Any primitive semisimple ring is isomorphic to* $\operatorname{End}_D(V)$ *for some division ring D and some finite dimensional vector space V over it.*

Any semisimple ring is isomorphic to a finite product of matrix rings over division rings. The numbers of factors, the isomorphism classes of the division rings, and the dimensions are uniquely determined.

Proof. Let R be a primitive semisimple ring, and let S be a faithful simple module. Writing R as S^n, we have:

$$R^{\circ} = \mathrm{End}_R(R) = \mathrm{End}_R(S^n) = \prod_{1 \le i \le n} \prod_{1 \le j \le n} \mathrm{Hom}_R(S, S) = M_n(E)$$

$M_n(E)$

where $E = \mathrm{End}_R(S)$ is a division ring by the last exercise, and $M_n(E)$ is the ring of $n \times n$ matrices over E (the last two equalities are by Exercise 2.3.15 and by the standard computation from linear algebra). We now note that the transpose is an isomorphism from $M_n(E)$ to $M_n(E^{\circ})$, so setting $D = E^{\circ}$ and $V = E^n$, we obtain the result.

The statement for general semisimple rings follows directly from this and Corollary 2.4.7. The uniqueness follows from Corollary 2.3.19 and Example 2.4.13 (applied with k a division ring). □

Remark 2.4.17. We could prove the result without passing through the Chinese Remainder Theorem, by applying the same calculation to the decomposition of a general semisimple ring. However, the Chinese Remainder Theorem and the notion of a primitive ring are important in their own right. □

Remark 2.4.18. Under the isomorphism from R to $M_n(D)$, the module D^n becomes isomorphic to S. This isomorphism depends on the choice of isomorphism from R to S^n as an R module. S is, by definition, a vector space over D, and the elments of R act on S by D-linear endomorphisms. Thus we have a *canonical* homomorphism of rings $R \to \mathrm{End}_D(S)$, that does not depend on any choices. Though we did not prove it directly, we now know that this is an isomorphism: The map is injective since R is primitive, and thus by dimension (over D) is also surjective. We shall return to this point later. The connection between the choices made is worked out in the following exercise. □

Exercise 2.4.19. Assume that a semisimple ring R has a unique simple module S. Choose an R-module isomorphism $f : S^n \to R = \mathrm{End}_R(R)^{\circ}$, and let S_i be the image of i-th component under f.

(1) Show that S_i is the set of endomorphisms with image in S_i
(2) Writing $1 = s_1 + \cdots + s_n$ with $s_i \in S_i$, show that $s_i s_j = \delta_{ij}$.
(3) Let $E = \mathrm{End}_R(S)$. Show that $\{f^{-1}(s_i)\}$ forms a basis of S over E.

End of lecture 6, Mar 30

We mention two additional corollaries of the structure theorem. The first explains why the notion of a semisimple ring is not very interesting in the commutative case:

Corollary 2.4.20. *A commutative ring is semisimple if and only if it is a finite product of fields.*

The second, rather surprising, is that we could consider right modules instead of left modules all along:

Corollary 2.4.21. *A ring R is semisimple if and only if R° is semisimple*

Before giving the proof, we introduce the notion of the dual module:

Exercise 2.4.22. Let V be a module over a ring D, and let $V^\vee = \mathrm{Hom}_D(V, D)$, viewed as an abelian group.

(1) If $\varphi \in V^\vee$ and $d \in D$, show that the map φd, defined as $(\varphi d)(v) = \varphi(v)d$, is in $V^\vee$
(2) Show that $(\varphi, d) \mapsto \varphi d$ gives $V^\vee$ the structure of a module over D° (this is called the *dual module*) **dual module**
(3) If $t : U \to V$ is a map of modules over D, show that composition with t determines a map $t^\vee$ of D° modules from $V^\vee$ to $U^\vee$, and that this process determined a contravariant functor from modules over D to modules over D°.
(4) Assume that D is a division ring. Show that applying duality twice is essentially the identity on finite-dimensional modules. More concretely, show that if U and V are finite dimensional, the map $t \mapsto t^\vee$ is a bijection.

Proof. If R is semisimple, it is isomorphic to a finite product of rings of the form $R_i = \mathrm{End}_{D_i}(V_i)$, with D_i a division ring, and V_i a finite dimensional vector space over it. The opposite is thus a product of the opposite rings ${R_i}^\circ$, i.e., to endomorphism algebras. The exercise shows that $\mathrm{End}_D(V)$ is canonically isomorphic to $\mathrm{End}^\circ_D(V^\vee)$, so is an endomorphism algebra itself. □

2.5. **Representations again.** In the case of the group algebra there is an additional ingredient we did not employ so far, namely, it is an algebra over a field. In this case, we can strengthen the result:

Proposition 2.5.1. *Let R be a semisimple k-algebra, where k is a field. Then every division ring D that appear in the decomposition of R contains k in its centre. If R is of finite dimension over k, then so is D, and if in addition k is algebraically closed, then $D = k$.*

Proof. The map from k to R corresponds to maps to each factor, so we reduce to the case $R = \mathrm{End}_D(V)$. Since k is in the centre of R, multiplication by elements of k is an endomorphism of the R-module structure, hence $k \subseteq D$, and for the same reason it is in the centre.

The second statement is obvious, since $D \subseteq R$. For the third, let $d \in D$. The sub-ring $k[d] \subseteq D$ is commutative (since k is in the centre), finite dimensional and has no zerodivisors, so is a finite field extension of k. Since k is algebraically closed, $d \in k$. □

Because of the proposition, we will mostly fix attention on the case that k is algebraically closed. In this case, our results can be summarised as follows:

Corollary 2.5.2. *Assume that G is a group of size n, and k is an algebraically closed field of characteristic p prime to n. Then:*

(1) *Each representation of G over k is a direct sum of irreducible representations. The multiplicity of each irreducible component is well defined.*
(2) *There are finitely many isomorphism classes of irreducible representations, $X_1, \dots, X_r$, of dimensions $n_1, \dots, n_r$. The multiplicity of X_i in the regular representation is n_i.*
(3) $\sum n_i^2 = n$

conjugacy class

class function

We next ask: can the number r of irreducible representations be described more directly in terms of the group? Recall that the group G acts on itself by conjugation, $c : G \times G \to G$ given by $c(g, h) = g^{-1}hg$. A *conjugacy class* in G is an orbit under this action. In other words, it is a set of the form $\{g^{-1}hg \,|\, g \in G\}$ for a fixed h. A function on G is called a *class function* if it is constant on conjugacy classes.

Proposition 2.5.3. *Let G be a group of size n, and k is an algebraically closed field of characteristic prime to n. Then the number of irreducible representation of G is equal to the number of conjugacy classes in G.*

Proof. We compute the centre of $k[G]$ in two different ways: First, an element $\sum_{g \in G} a_g g \in k[G]$ is in the centre if and only if it commutes with all elements of G, i.e., for all $h \in G$

$$\sum a_g g = h^{-1} \sum a_g gh = \sum a_g h^{-1}gh = \sum a_{hgh^{-1}} g$$

Since the elements of G are linearly independent in $k[G]$, we get that this happens if and only if $g \mapsto a_g$ is a class function.

On the other hand, we have $k[G] = M_1 \times \ldots \times M_r$, where r is the number of irreducible representations, and each M_i is a matrix algebra over k. Since the centre of each M_i is k (the scalar matrices), we have $Z(k[G]) = k^r$.

It follows that the dimension of the space of class functions is r, hence that there are r conjugacy classes. □

Exercise 2.5.4. Compute the conjugacy classes and the irreducible representations of S_3

End of lecture 7, Apr 5

The last result is somewhat unsatisfactory: though we found that conjugacy classes and irreducible representations are the same in number, we did not find an explicit matching between them. The last exercise also does not provide an obvious way to match them. We now reprove the last result in a manner that explains the situation. We refer to [Ser77] for details on this, and for more information about representations.

Unless otherwise stated, we now assume that *k is an algebraically closed field of characteristic* 0 (for example, the complex numbers). We first note, that it is unlikely to expect that a canonical matching between conjugacy classes and irreducible representations be found. The reason is that, while the structure of representations depends only on the group algebra, the passage from the group to the group algebra forgets information. For example, according to Corollary 2.4.20, any two commutative groups have isomorphic group algebras. Instead, we will find a canonical isomorphism between the linear space spanned by the irreducible representations and the linear space of class functions. In fact, we can do it on the level of an arbitrary semisimple k-algebra.

We begin by noting that there are two canonical representations associated with G: the group algebra $R = k[G]$ (viewed as a module over itself), and the regular representation k^G. What is the relation between them? Any group element determines, by evaluation a linear functional on k^G. Extending by linearity, we obtain an isomorphism of vector spaces between the regular representation and $R^\vee$, the vector space dual of the group algebra. The representation itself is obtained by considering the *right* action of R on itself.

Inside the group algebra R we have the centre $Z(R)$. On the other hand, inside $R^\vee$ we have (under the above isomorphism) the class functions. Can we describe these functions in terms of the algebra structure?

Exercise 2.5.5. Let $R = k[G]$ and let $\varphi : R \to k$ be a linear function. Show that φ is a class function if and only if $\varphi(rs) = \varphi(sr)$ for all $r, s \in R$.

We may now forget the group entirely, and describe the situation on the level of algebras. An element $\varphi \in R^\vee$ satsfies $\varphi(rs) = \varphi(sr)$ for all $r, s \in R$ precisely if it is 0 on the subspace $[R, R]$ of R spanned by the elements $[r, s] = rs - sr$. Thus, the class functions can be described as the subspace $(R/[R, R])^\vee$ of $R^\vee$. Note that $R/[R, R]$ is the largest commutative quotient of R.

The point now is that the two sides are canonically isomorphic:

Proposition 2.5.6. *Let R be a semisimple algebra over an algebraically closed field k of characteristic* 0. *For each $r \in R$, let $T_r : R \to R$ be the right multiplication by r (so $T_r(s) = sr$), and let $d_r : R \to k$ be given by $d_r(s) = \mathrm{Tr}(T_{sr})$. Then:*

1. *$r \mapsto d_r$ is an isomorphism $d : R \to R^\vee$ of R-modules.*
2. *The restriction of d to the centre $Z(R)$ of R identifies it with the subspace $(R/[R, R])^\vee$ of $R^\vee$.*

The following diagram describes the situation:

$$\begin{tikzcd}
R \arrow[r, "\sim", "d"'] & R^\vee \\
Z(R) \arrow[u, hook] \arrow[r, "\sim", "d"'] & (R/[R,R])^\vee \arrow[u, hook]
\end{tikzcd} \tag{2.1}$$

The statement will follow from the following analogous statement in linear algebra:

Lemma 2.5.7. *Let U be a finite dimension vector space over a field k, and let $V = End_k(U)$*

(1) The map $T \mapsto (S \mapsto \mathrm{Tr}(TS))$ determines an isomorphism from V to $V^\vee$
(2) The space $[V, V]$ spanned by the commutators in V coincides with the elements of trace 0

Proof. (1) Will be proved in §3.2.
(2) Since $\mathrm{Tr}(ST) = \mathrm{Tr}(TS)$ for all $T, S \in V$, we need to show that any trace-zero transformation is a linear combination of commutators. Picking a basis, we may identify V with matrices, and the space of matrices of trace 0 is spanned by matrices from the standard basis off the main diagonal, and diagonal matrices with two non-zero entries equal to 1 and -1. This reduces to the two dimensional case, and in fact to the cases $e = \left(\begin{smallmatrix} 0 & 1 \\ 0 & 0 \end{smallmatrix}\right)$ and $h = \left(\begin{smallmatrix} 1 & 0 \\ 0 & -1 \end{smallmatrix}\right)$. We have $e = [f, e]$ and $h = [e, e^t]$, where $f = \left(\begin{smallmatrix} 1 & 0 \\ 0 & 0 \end{smallmatrix}\right)$, and e^t is the transpose of e. □

The proof of the proposition reduces to this via Theorem 2.4.16:

Proof of 2.5.6.
(1) If $t, s \in R = \prod \mathrm{End}_k(U_i)$, we may write $t = (t_1, \dots, t_r)$, $s = (s_1, \dots, s_r)$, with $t_i, s_i \in \mathrm{End}_k(U_i)$. Then $st = \sum s_i t_i$, and since the trace is additive with respect to direct sums, we have $d_t(s) = \sum d_{t_i}(s_i)$. Hence, we may assume that $t = t_i \neq 0$, and we need to find $s \in \mathrm{End}_k(U_i)$ with $\mathrm{Tr}(T_{st}) \neq 0$, where T_{st} is viewed as a map on $\mathrm{End}(U_i)$. But $\mathrm{End}(U_i)$ is isomorphic to the direct sum of n_i copies of U_i, with st acting diagonally, so $\mathrm{Tr}(T_{st})$ is n_i times the trace of st acting on U_i. Since we are in characteristic 0, the lemma shows this is non-degenerate.
(2) If c is in the centre, and $s, t \in R$, we have $\mathrm{Tr}([s, t]c) = \mathrm{Tr}(stc - tsc) = \mathrm{Tr}(stc - tcs) = \mathrm{Tr}([s, tc]) = 0$. Thus $d_c \in (R/[R, R])^\vee$. For the converse, if $d_c(t) = 0$ for all $t \in [R, R]$, reducing again to the case of matrices and using the second part of the lemma, we see that $\mathrm{Tr}(tc) = 0$ whenever $\mathrm{Tr}(t) = 0$. It follows that there is a scalar d such that $\mathrm{Tr}(t(c - d)) = 0$ for all t. Hence $c - d = 0$, so c is multiplication by d. □

Exercise 2.5.8. Show that if $R = k[G]$ is a group algebra of dimension n, then d sends $g \in G$ to the function $n\delta_{g^{-1}} \in k^G$ (where δ_g is the characteristic function of g)

Let I be the set of isomorphism classes of simple modules for R. Thus, I has size r, and to each element $s \in I$ corresponds an idempotent $e_s \in Z(R)$ with $1 = \sum_{s \in I} e_s$. The e_s form a basis for $Z(R)$ over k, and so we have recovered Proposition 2.5.3 and upgraded it to the isomorphism d.

We conclude with a number of observations.

Exercise 2.5.9. Let S be a simple R-module of dimension n over k. We write $[S] = \frac{1}{n} e_S \in Z(R)$. If X is an arbitrary finite-dimensional R-module, we may write $X = \oplus_{S \in I} S^{n_s}$ for unique numbers n_s, and we set $[X] = \sum n_s[S]$. Thus we have a well defined map from isomorphism classes of finite-dimensional R-modules to $Z(R)$.

On the other hand, to any such X we may assign an element t_X of $(R/[R,R])^\vee$, given by $t_X(r) = \mathrm{Tr}(r_X)$, where r_X is the linear map on X given by r.

Show that these two maps are compatible with d: $d([X]) = t_X$. In the language of representation theory, t_X is called the *character* of X. Show that if $t_X = t_Y$ then X and Y are isomorphic.

character

Exercise 2.5.10. The isomorphism d determines (and is determined by) the symmetric bilinear map $\langle \cdot, \cdot \rangle$ on $R^\vee$ (and on $(R/[R,R])^\vee$) given by $\langle f, g \rangle = g(d^{-1}(f))$.

(1) Show that the maps of the form t_X, where X is a simple module, form an *orthonormal* basis of $(R/[R,R])^\vee$ for this bilinear map. Conclude that $\langle t_Y, t_Y \rangle = 1$ if and only if Y is simple.

Assume now that G is a group of size n, and $R = k[G]$. We denote by $G /\!\!/ G$ the set of conjugacy classes, and for $c \in G /\!\!/ G$, by $f_c \in k^{G/\!\!/G} = (R/[R,R])^\vee$ the characteristic function of c. Note that if g and h are cojugate, then so are g^{-1} and h^{-1}, so the inverse c^{-1} of a conjugacy class c is well defined.

$G /\!\!/ G$

(2) Show that $\langle t, s \rangle = \frac{1}{n} \sum_{g \in G} t(g)s(g^{-1})$ on k^G.

(3) Show that if t is a class function, then $\langle t, f_c \rangle = \frac{|c|}{n} t(c^{-1})$. In particular, the transition matrix between the two bases of the space of class functions is given by $(\frac{|c|}{n} t_X(c^{-1}))$, where c ranges over the conjugacy classes, and X over isomorphism classes of irreducible representations.

(4) Show that for $g, h \in G$ we have

$$\sum_{X \in I} t_X(g) t_X(h^{-1}) = \begin{cases} n/|c| & g, h \in c \text{ for some class } c, \\ 0 & \text{otherwise.} \end{cases}$$

(Hint: expand f_c in the basis $\{t_X\}$)

The transition matrix between the two canonical bases, computed above, is called the *character table* of G. It contains all the information about the representations of G.

character table

Exercise 2.5.11. Let G be a finite group, and let $G^\vee$ be the set of one-dimensional representations of G, viewed as group homomorphisms into k^*. The pointwise product gives $G^\vee$ the structure of an abelian group, called the *Pontrjagin dual group* of G.

Pontrjagin dual group

(1) Assume that G is abelian. Show that $G^\vee$ forms a basis of k^G. Hence k^G can be canonically identified (as a k-algebra) with $k[G^\vee]$. The transition matrix (character table) discussed above is called the *Fourier transform* in this case.

Fourier transform

(2) Show that in the case $G = \mathbb{Z}/n\mathbb{Z}$, $G^\vee = \mu_n$, the group of n-th roots of 1. Compute the transition matrix (character table) in this case.

(3) If G is an arbitrary finite group, show that $k[G^\vee]$ is isomorphic to $k[G^{ab}]$

This completes our discussion of representations over an algebraically closed field of characteristic 0. We will consider the non-algebraically closed case in the next subsection. Now we consider the case when the characteristic p of k divides n. We had seen that in this case semisimplicity fails, We will now see that in the extreme case that G is a p-group, it fails in the strongest way possible. We recall that every p-group (i.e., a group whose order is a power of p) has a non-trivial centre $Z(G)$ (this follows from counting the sizes of conjugacy classes).

Proposition 2.5.12. *If V is a finite-dimensional non-trivial representation of a p-group G over an algebraically closed field k of characteristic p, then V has a non-zero fixed vector. In particular, every irrducible representation of G is trivial (hence* 1*-dimensional)*

Proof. We first claim that every 1-dimensional representation of G over k is trivial. Indeed, such a representation is a group homomorphism into the multiplicative group of k, and the order of each element in the image is a power of p. The only such element in a field of characteristic p is 1.

Assume that G has a proper non-trivial normal subgroup H. By induction, the space $U = V^H$ of vectors fixed by H is non-zero, and is a representation of G/H. Again by induction, U has a vector fixed by G/H, hence by G.

It remains to deal with the base of the induction, i.e., when G is simple. Then G must be abelian, since otherwise the centre is a non-trivial proper normal subgroup. Hence G is cyclic, generated by one element g, which determines a linear automorphism of V, and since k is algebraically closed, it has an eigenvector v. This eigenvector then spans a one-dimensional representation, which must be trivial by the above, hence v is fixed. □

End of lecture 8, Apr 19

3. Tensor products

In this section, we would like to answer the following questions:

(1) How to prove Lemma 2.5.7(1)?
(2) Suppose that S is an R-ring, and M is an R module. Is there a way to turn M into an S-module in a canonical way?
(3) If G is a group, the space of class functions forms a sub-algebra of k^G. We have seen that this space is spanned by characters of representations. Can the product be described in terms of representations as well?
(4) What is the coproduct in the category of commutative rings? How about general rings?
(5) What can be said of representations over non-algebraically closed fields?

The common motive for all these questions is that they involve a kind of a product between elements of abelian groups (or modules). By that we mean a bilinear map $b : U \times V \to W$, where U, V and W are abelian groups. We may think of $b(u, v)$ as a "product" of u and v, with value in W, and bilinearity means that this product is distributive. Fixing U and V, this still may seem potentially complicated, as we are dealing with arbitrary W, and with bilinear, rather than linear maps. Fortunately, there is a universal one:

Definition 3.0.1. Let U and V be two abelian groups. The *tensor product* of U and V consists of an abelian group $U \otimes V$ and a bilinear map $\otimes : U \times V \to U \otimes V$, $\langle u, v\rangle \mapsto u \otimes v$, which is universal: if $b : U \times V \to W$ is a bilinear map, there is a unique homomorphism $b' : U \otimes V \to W$ with $b'(u \otimes v) = b(u, v)$ for all $u \in U, v \in V$.

tensor product

$U \otimes V$

As usual with universal objects, the tensor product is uniquely determined by its property (up to a unique isomorphism).

Exercise 3.0.2. Let U and V be abelian groups. Show the following:

1. $U \otimes V$ is canonically isomorphic to $V \otimes U$, and if W is another abelian group, $(U \otimes V) \otimes W$ is isomorphic to $U \otimes (V \otimes W)$.
2. $\mathbb{Z} \otimes U$ is isomorphic to U. If $u \in U$ and $v \in V$ have coprime orders, then $u \otimes v = 0$ (in particular, the tensor product of two groups of coprime orders is 0).
3. For fixed U, $V \mapsto U \otimes V$ is a functor: an additive map $f : V \to W$ induces a map $1 \otimes f : U \otimes V \to U \otimes W$, compatibly with compositions.

Example 3.0.3. If A is an abelian group, a ring structure on A is a bilinear map $\cdot : A \times A \to A$, hence is given by a linear map from $A \otimes A$ to A. Conversely, such a linear map determines a (non-unital) ring structure if it is associative. □

We did not yet state that the tensor product exists. Since $U \otimes V$ is the "most free" product of U and V, it is simplest to construct it freely:

Exercise 3.0.4. Show that the tensor product of any two abelian groups U and V exists. Show that if U is generated by $\{u_\alpha\}$ and V is generated by $\{v_\beta\}$, then $U \otimes V$ is generated by $\{u_\alpha \otimes v_\beta\}$ (consider the free abelian group on the set $U \times V$, and divide by suitable relations)

3.1. **Ring structure.**

Assume now U is a module over a ring R. By Exercise 3.0.2, we have an induced R-module structure on $U \otimes V$, for all abelian groups V. This R-module structure is uniquely determined by the following property: If W is an R-module, there is a one-to-one correspondence between R-module homomorphisms $U \otimes V \to W$ and bilinear maps $b : U \times V \to W$ satisfying $b(ru, v) = rb(u, v)$ for all $r \in R$, $u \in U$ and $v \in V$. Thus, it is a form of associativity.

Assume now that both U and V are R-modules. It is then natural to be interested in bilinear maps on $U \times V$ in which multiplication by elements of R on each of the coordinates gives the same result. This happens, for example, for inner products, or for algebra structures. One easily verifies that this requirement is reasonable only if U and V are modules on opposite sides, i.e., one of them is a module over R°. As with usual bilinear maps, it is convenient to consider the universal one:

Definition 3.1.1. Let R be a ring, V an R-module, and U an R°-module. An *R-bilinear map* on $U \times V$ is a bilinear map $b : U \times V \to W$, where W is an abelian group, satisfying $b(ur, v) = v(u, rv)$ for all $r \in R$ and $(u, v) \in U \times V$.

R-bilinear map

The *tensor product over R* of U and V, denoted $U \otimes_R V$, is a universal R-bilinear map from $U \times V$.

tensor product over R

$U \otimes_R V$

Thus, we have $ur \otimes v = u \otimes rv$ in $U \otimes_R V$. Most of Exercise 3.0.2 can be repeated in this context:

Exercise 3.1.2. Let U be an R°-module. Show:

(1) $U \otimes_R R = U$ (canonical isomorphism; the left side is apriori only an abelian group, but gains a right R-module structure with the action of R on itself on the right, and then this is true as R°-modules)

(2) If V is an R-module, $u \in U$ and $v \in V$, and $1 \in \operatorname{ann} u + \operatorname{ann} v$, then $u \otimes v = 0$ in $U \otimes_R V$[4] (Note that $\operatorname{ann} u + \operatorname{ann} v$ is, in general, only an additive subgroup of R.)

(3) The assignment $V \mapsto U \otimes_R V$ is a functor from R-modules to abelian groups.

(4) The usual tensor product is the same as tensor product over $\mathbb{Z}$

The existence of tensor products over R follows formally from that of abelian groups:

Proposition 3.1.3. *For any ring R, R-module V and R°-module U, the tensor product $U \otimes_R V$ exists.*

Proof. Set $U \otimes_R V = U \otimes V / \{ur \otimes v - u \otimes rv \mid r \in R, u \in U, v \in V\}$. The map sending (u, v) to the class of $u \otimes v$ is clearly R-bilinear, and it is easy to see that it is universal. □

We stress that in general, $U \otimes_R V$ has only the structure of an abelian group: there is no well defined R-module structure on it. However, if U is, in addition, and S-module for some ring S, and the two module structures on U commute (i.e., $(su)r = s(ur)$ for $s \in S$, $r \in R$ and $u \in U$), then $U \otimes_R V$ becomes an S-module as above (by Exercise 3.1.2). An important special case is when R is commutative, and $S = R$, with the same module structure. Thus, for commutative rings R, $U \otimes_R V$ is canonically an R-module, and we

[4] If R is the (commutative) ring of functions on some space X, we may think of u and v as generalised functions on X. Then u is supported on the set of common zeroes of elements of $\operatorname{ann} u$ (and similarly for v), and the condition $1 \in \operatorname{ann} u + \operatorname{ann} v$ means that the supports are disjoint. Thus, the geometric content here is that the product of functions supported on disjoint subsets is 0.

treat it as such (note that the R-module structure is the same if we use the R-module structure on V instead).

Exercise 3.1.4. Let k be a field, and let U and V be k-vector-spaces. Show that if $\{u_\alpha\}$ and $\{v_\beta\}$ are bases for U and V, respectively, then $\{u_\alpha \otimes v_\beta\}$ is a basis for $U \otimes_k V$. In particular, $\dim_k(U \otimes_k V) = \dim_k(U)\dim_k(V)$.

Exercise 3.1.5. Show that if F is a free left R-module, and M is a right R-module, then $M \otimes_R F$ is isomorphic to a direct sum of copies of M. Show that if $N \subseteq M$ is a submodule, then $N \otimes_R F$ is a submodule of $M \otimes_R F$ (we say that F is a *flat R-module*, so every free module is flat). In particular, this holds for any module over a division ring

flat R-module

3.2. **Some linear algebra.** We now fix a field k. All our tensor products will be over k, so we suppress the subscript. Let U and V be k-vector spaces. If $\varphi \in U^\vee$ and $v \in V$, the map $u \mapsto \varphi(u)v$ is a linear map from U to V. Hence, we have a map from $U^\vee \times V$ to $\operatorname{Hom}(U, V)$, which is clearly bilinear.

Proposition 3.2.1. *Let U and V be k-vector spaces.*

1. *The linear map $U^\vee \otimes V \to \operatorname{Hom}(U, V)$ is injective. If U or V are finite dimensional, the map is an isomorphism*
2. *The linear map $U^\vee \otimes V^\vee \to (U \otimes V)^\vee$ induced by $(\varphi, \psi) \mapsto (u \otimes v \mapsto \varphi(u)\psi(v))$ is injective, and is bijective if either U or V is finite-dimensional.*
3. *If W is another space, the composition map $\operatorname{Hom}(U, V) \otimes \operatorname{Hom}(V, W) \to \operatorname{Hom}(U, W)$ restricts to the map determined by $(\varphi \otimes v, \psi \otimes w) \to \psi(v)\varphi \otimes w$.*

Proof. By Exercise 3.1.4, we may write an arbitrary element of $U^\vee \otimes V$ as $\sum \varphi_i \otimes v_i$, with each of $\{\varphi_i\}$ and $\{v_i\}$ independent. Hence, if $u \in U$ is any element such that not all $\varphi_i(u) = 0$, the corresponding linear map will not map it to 0.

If U is finite dimensional, let $\{u_i\}$ be a basis. An arbitrary linear map T from U is determined by its values $v_i = Tu_i$ on the basis. Then T is represented by $\sum \varphi_i \otimes v_i$, where φ_i is the basis dual to u_i. The proof in the other case is similar.

The other parts are similar and left as an exercise □

Exercise 3.2.2. Complete the proof of the proposition.

Let U be a finite dimensional vector space. The map from $U^\vee \times U$ to k given by $(\varphi, u) \mapsto \varphi(u)$ is bilinear, hence induces a linear *evaluation* map $\operatorname{ev} : U^\vee \otimes U \to k$. It is natural to ask how to interpret this map if we identify $U^\vee \otimes U$ with $\operatorname{End}(U)$, as above.

Proposition 3.2.3. *Let U be a finite dimensional space. The evaluation map* $\operatorname{ev} : U^\vee \otimes U \to k$ *coincides with the trace map under the isomorphism from Prop. 3.2.1*(1)

Proof. Since both maps are linear, it suffices to prove they coincide on a basis. If u_i is a basis for U and φ_i is the dual basis, a basis for $U^\vee \otimes U$ is given by $\varphi_i \otimes u_j$ (by Exercise 3.1.4). On this basis the claim is obvious. □

Proof of Lemma 2.5.7(1). We identify $\operatorname{End}(U)$ with $U^\vee \otimes U$. We first note that according to Prop. 3.2.1(2) we have isomorphisms

$$U^\vee \otimes U \xrightarrow{\sim} U \otimes U^\vee \xrightarrow{\sim} (U^\vee)^\vee \otimes U^\vee \xrightarrow{\sim} (U^\vee \otimes U)^\vee$$

Tracing through, this isomorphism sends $\varphi \otimes u \in U^\vee \otimes U$ to the map sending $\psi \otimes v$ to $\psi(u)\varphi(v)$. On the other hand, the map described in the lemma is given by $\varphi \otimes u \mapsto (\psi \otimes v \mapsto \operatorname{Tr}(\varphi \otimes u \circ \psi \otimes v))$. We have $\varphi \otimes u \circ \psi \otimes v = \psi(u)\varphi \otimes v$ by Prop. 3.2.1(3), and its trace is $\psi(u)\varphi(v)$ by Prop. 3.2.3. So the two maps coincide. □

End of lecture 9, Apr 20

3.3. **Tensor product representations.** Assume now that U and V are representations of a group G. Then $U \otimes V$ admits the structure of a G-representation, given by the action on each factor. A calculation similar to the above shows:

Exercise 3.3.1. If T and S are endomorphisms of finite-dimensional vector spaces U and V (respectively), and $T \otimes S$ is the resulting endomorphism of $U \otimes V$, then $\operatorname{Tr}(T \otimes S) = \operatorname{Tr}(T)\operatorname{Tr}(S)$

Applying this to each element of G separately, we obtain:

Corollary 3.3.2. *If X and Y are (finite dimensional) representations of a group G, then $t_{X\otimes Y} = t_X t_Y$ (as functions from G to k)*

We had seen that for finite groups G and algebraically closed fields k (of suitable characteristic), the character of a representation determines its isomorphism type, so this is helpful in computations.

Example 3.3.3. The irreducible representations of S_3 are the trivial representations $\mathbf{1}$, the sign representation S, and a 2-dimensional representation X. Denoting the conjugacy classes of S_3 by $1, 2, 3$, according to their orders, a direct computation shows: $t_X(1) = 2$, $t_X(2) = 0$ and $t_X(3) = -1$. It follows from the last corollary that $t_{X\otimes X} = (4, 0, 1) = t_X + t_S + t_{\mathbf{1}}$, so $X \otimes X$ is isomorphic to $X \oplus S \oplus \mathbf{1}$ □

To appreciate the power of the method from the last example, we would like to describe the isomorphism directly. As a first step, we consider dual representations:

Exercise 3.3.4. Let U and V be representations of a group G over a field k. For $T \in \operatorname{Hom}_k(U, V)$ and $g \in G$ we let $g(T) \in \operatorname{Hom}_k(U, V)$ be the map $g(T)(u) = gT(g^{-1}u)$. Show the following.

1. $T \mapsto g(T)$ is a representation of G. What are the fixed vectors?
2. With this structure, the maps in Prop. 3.2.1 are maps of representations (note that $U^\vee$ is a representation as a special case of the definition, with V the trivial representation)

(3) In the context of §2.5, show that $t_U^\vee(g) = t_U(g^{-1})$. Conclude that if in G, each element is conjugate to its inverse, then each representation is isomorphic to its dual (note that this holds for any S_n).
(4) Assume that $k = \mathbb{C}$, the complex numbers, then $t_U^\vee(g) = t_U\bar{(}g)$, where $\bar{x}$ is the complex conjugate of x. Conclude that if t_U is real-valued, then U is isomorphic to its dual (again, this is the case for S_n) *Hint*: consider the eigenvalues of an element g acting on U, and use the previous part.
(5) Generalize the previous part to an arbitrary field (notice that the characters take values in an extension of the prime field by sufficiently many roots of 1, and examine the Galois group to find a replacement for complex conjugation)
(6) Recall that the representation X from the example can be realised as the subspace $\{(x_1, x_2, x_3) \,|\, x_1 + x_2 + x_3 = 0\}$ of the standard representation $V = k^3$. Show that $X^\vee$ can be identified with the subspace of $V^\vee$ consisting of functionals killing the constant vectors. Find directly an isomorphism from $X^\vee$ to X in this realisation.

Example 3.3.5. We continue with Example 3.3.3 and its notation. Our goal is to describe explicitly the isomorphism from $\mathbf{1} \oplus S \oplus X$ to $X \otimes X$. The last exercise shows (in three different ways) that X is isomorphic to $X^\vee$, and that $X^\vee \otimes X$ is isomorphic to $\mathrm{Hom}_k(X, X) = \mathrm{End}_k(X)$, so we need to find injective maps from each of $\mathbf{1}, S, X$ to $\mathrm{End}_k(X)$ (the images will be automatically independent, since these are different irreducible representations).

A map from $\mathbf{1}$ corresponds to a fixed vector, i.e., to and endomorphism of X as a representation. One such choice is the identity, $\mathbf{1}$ corresponds to the scalar endomorphisms.

Likewise, the subspace corresponding to S is the set of endomorphisms T satisfying $T(gv) = gT(v)$ for for even g, but $T(gv) = -gT(v)$ for transpositions. To find such a T, write $X = X_1 \oplus X_2$, where X_1 and X_2 are the two eigenspaces for the action of the even elements (explicitly, each is spanned by a vector $(1, \alpha, \alpha^2)$, where α is a third root of 1). It is then clear that each T with the same eigenspaces will commute with the even elements. To satisfy the requirement for transpositions, T needs in addition to have inverse eigenvalues (i.e., trace 0).

Finally, each of the spaces X_i maps to the space of endomorphisms T that map X_{3-i} to X_i, and X_i to 0 (i.e., to strict upper and strict lower triangular matrices, in a basis of eigenvectors for even elements). □

3.4. **Base change.** Assume now that S is a ring over R, and that M is an R-module. The R-ring structure on S makes it a right R-module, and we may thus use Definition 3.1.1 to form the tensor product $S \otimes_R M$. The ring structure on S exhibits each $s \in S$ as an R°-module map $a \mapsto sa$, and so by functoriality determines an S-module structure on $S \otimes_R M$. The resulting module, known as the *base change* (or *extension of scalars*) of M to S, is the **base change** **extension of scalars** "most free" extension of M to a module over S. More precisely, it satisfies the following property.

Proposition 3.4.1. *Given an S module N, and an R-module homomorphism $f : M \to N$, there is a unique S-module map $\tilde{f} : S \otimes_R M \to N$ such that $\tilde{f}(1 \otimes m) = f(m)$ for all $m \in M$. In other words, $S \otimes_R M$ is universal among maps from M to an S-module.*

Exercise 3.4.2. (1) Prove Proposition 3.4.1

(2) Show that if M is given in terms of generators and relations over R, then $S \otimes_R M$ is given by the same generators and relations over S.

As an example, we consider representations again. It may be easier to understand representations of a group G in terms of representation of subgroups of H. Since a representation is just a module over the group algebra, we may obtain a G-representation from an H-representation X via the base change $k[G] \otimes_{k[H]} X$. This is called the *induced representation* from H to G of the representation X, denoted $\operatorname{Ind}_H^G(X)$.

induced representation

The defining property of the base change implies that there is a bijection from the set $\operatorname{Hom}_G(\operatorname{Ind}_H^G(X), Y)$ to the set $\operatorname{Hom}_H(X, Y)$ for any G-representation Y. In fact, this is clearly a k-linear isomorphism between the two k-vector spaces. In particular, they have the same dimension. When Y is an irreducible representation, this dimension is (by Schur's Lemma and semisimplicity) nothing but the coefficient of t_Y in the expansion of $\operatorname{Ind}_H^G(X)$ in terms of irreducible representations. In other words, in the terminology of Exercise 2.5.10, we proved:

Proposition 3.4.3 (Frobenius reciprocity theorem)**.** *If H is a subgroup of a finite group G, then for any representation X of H and Y of G we have $\langle \operatorname{Ind}_H^G(X), Y\rangle_G = \langle X, Y\rangle_H$ (where $\langle\cdot,\cdot\rangle_K$ is the bilinear map for the group K, cf. Exercise 2.5.10)*

Example 3.4.4. Let H be the subgroup of order 3 of S_3. We had seen that the irreducible representations of H are one-dimensional, of the form V_α where $\alpha : H \to \mu_3$ is a group homomorphism to the group μ_3 of 3-rd roots of 1. Let $X_\alpha = \operatorname{Ind}_H^{S_3} V_\alpha$ be the induced representation. Then we have for an irreducible representation U of S_3 (with the notations of Example 3.3.3)

$$\langle X_\alpha, U\rangle_{S_3} = \langle V_\alpha, U\rangle_H = \begin{cases} 1 & \alpha = 1, U = \mathbf{1} \text{ or } U = S \\ 1 & \alpha \neq 1, U = X \\ 0 & \text{otherwise} \end{cases}$$

Hence, $X_\alpha = \mathbf{1} \oplus S$ if $\alpha = 1$, and $X_\alpha = X$ if α is non-trivial. □

Exercise 3.4.5. Compute the representations of S_3 induced from irreducible representations of subgroups of order 2.

3.5. **Tensor products of rings.** If S and T are rings, we may form $S \otimes T$ as an S-module (or a T°-module). However, it also has a ring structure, given by $(s_1 \otimes t_1) \cdot (s_2 \otimes t_2) = s_1 s_2 \otimes t_1 t_2$. The rings S and T each embed into $S \otimes T$, and their images commute. In fact, it is universal for this property:

Proposition 3.5.1. *Given rings S, T and R, and ring homomorphisms $f : S \to R$ and $g : T \to R$ such that $f(s)g(t) = g(t)f(s)$ for all $s \in S$ and $t \in T$, there is a unique map of rings $h : S \otimes T \to R$ with $h(s \otimes 1) = f(s)$ and $h(1 \otimes t) = g(t)$ for all $s \in S$ and $t \in T$. In particular, $\otimes$ is the coproduct in the category of commutative rings.*

Exercise 3.5.2. Prove the proposition

We would like define, in a similar way, a ring structure on $S \otimes_R T$, where S and T are rings over R. However, this is in general ill defined: for example, in $R \otimes_R R$ we would have:

$$1 \otimes r_1 r_2 = r_1 \otimes r_2 = (r_1 \otimes 1)(1 \otimes r_2) = (1 \otimes r_2)(r_1 \otimes 1) = \\ (r2 \otimes 1)(r_1 \otimes 1) = r_2 r_1 \otimes 1 = 1 \otimes r_2 r_1 \quad (3.1)$$

for any $r_1, r_2 \in R$, which is absurd if R not commutative. But if R is commutative, and further lies in the centre of S and T (i.e., when S and T are algebras over R), then $S \otimes_R T$ has a ring structure as before. Furthermore, there is then a canonical R-algebra structure on $S \otimes_R T$, and we get:

Proposition 3.5.3. *If S and T are algebras over a (commutative) ring R, then so is $S \otimes_R T$. The maps $S \to S \otimes_R T$ and $T \to S \otimes_R T$ are R-algebra maps with commuting images, and $S \otimes_R T$ is universal with this property.*

In particular, $\otimes_R$ is the coproduct in the category of commutative R-algebras.

Again, the proof is left as an exercise.

We still did not describe the coproduct in the category of all rings (or R-algebras). The description is straightforward but a bit messy: it is obtained as a union of suitable quotients of the abelian groups (or R-modules) $S \otimes T$ tensored with itself n times, as n goes to infinity (this is completely analogous to the construction of the free product of monoids, which is the coproduct in that category).

End of lecture 10, Apr 26

3.6. **Field extensions.** Lastly, we consider tensor products of finite field extensions. Let K be a field, and let E be a fixed algebraically closed field extending K. We recall the following definitions:

Definition 3.6.1. Let L be a finite extension of K of degree n. L is a *normal extension* if all embeddings of L in E over K have the same image. It is a *separable extension* if the number of such embeddings is n. It is a *Galois extension* if it is both normal and separable.

normal extension

separable extension

Galois extension

It is easy to show that these notions do not depend on the chosen algebraically closed field E. We fix it E till the end of the section.

Let L be a fixed finite extension, and denote by X the set of embeddings of L in E over K. Let $G = Aut(L/K)$ be the group of automorphisms of L over K. Then G acts of X by composition, without fixed points. If $a, b \in X$ have the same image, then $a^{-1} \circ b$ is a well-defined element of G, which takes a to b.

On the other hand, let A be any E-algebra. Any algebra map from A to E can, of course, be regarded as an element of $A^\vee$, the vector space dual of A over E. We need the following observation:

Lemma 3.6.2. *Let A be a D-ring, where D is a division ring. Then any collection of D-ring homomorphisms $A \to D$ is linearly independent as a subset of $A^\vee$.*

Proof. We first claim that for any two distinct D-ring maps $f, g : A \to D$ there is $a \in A$ with $f(a) = 0$ and $g(a) = 1$: since f and g are distinct, we have $b \in A$ with $c := f(b) \neq g(b)$. Then $f(b - c) = f(b) - f(c) = c - c = 0$ and $d := g(b - c) = g(b) - c \neq 0$. So $a = d^{-1}(b - c)$ solves the problem.

Now assume that $\sum_{i=0}^{k} f_i d_i = 0$ for some D-ring maps $f_i : A \to D$ and $d_i \in D$. For $1 \leq i \leq k$, let $a_i \in A$ be such that $f_0(a_i) = 1$ and $f_i(a_i) = 0$, and put $a = \prod a_i$. Then $f_0(a) = 1$ and $f_i(a) = 0$ for $i > 0$. Plugging this into the linear dependence, we find $d_0 = 0$. The same argument shows that $d_i = 0$ for each i. □

Remark 3.6.3. If A was the ring of D-valued functions on some set S (with D embedded in A as the constant functions), each point x of S determines a D-ring map $p_x : A \to D$, given by evaluation at x. The value of the function $a \in A$ at x is thus given by $p_x(a)$.

Hence, we may view the set of all D-ring maps from a D-ring A to D as the set of points of some space S, with A the algebra of D-valued functions on it. If we give D the topology in which the finite subsets (and D itself) are closed, the weak topology on S making all functions $a \in A$ continuous is called the *Zariski topology*. The last lemma then asserts that this topology is T_1 (points are closed).

Zariski topology

Note that the set X of embeddings of L in E is precisely of this form. □

Combining the observations so far, we may reformulate the definitions as follows:

Corollary 3.6.4. *Let X be the set of embeddings of a finite field extension L of K in E, and let $G = \mathrm{Aut}(L/K)$. Then:*

1. *L/K is normal if and only if the action of G on X is simple and transitive (this means that the combined map $G \times X \to X \times X$ given by the projection and the action is an isomorphism)*
2. *L/K is separable if and only if X forms a basis for $(E \otimes_K L)^\vee$*

We would like to reformulate this condition in a language that does not mention E. If $g \in G$, together with the identity map we obtain from the coproduct condition a map of rings $m_g : L \otimes_K L \to L$ (given explicitly by $m_g(x \otimes y) = xg(y)$). Combining these maps for all $g \in G$, we get a map $m : L \otimes_K L \to L^G$ (given by $m(x)(g) = m_g(x)$).

Recall that a *nilpotent element* of a ring A is an element $a \in A$ with $a^n = 0$ for some n. A is a *reduced* ring if the only nilpotent element in it is 0. We will need the following fact, which will be proved later.

nilpotent element

reduced

Proposition 3.6.5. *Let A be a finite-dimensional algebra over a field k. Then A is semisimple if and only if it is reduced*

We note that by Corollary 2.4.20, a reduced commutative finite-dimensional k-algebra is a finite product of finite field extensions of k.

Proposition 3.6.6. *Let L be a finite extension field of K. Then*

(1) L/K is separable if and only if $L \otimes_K L$ is reduced

(2) L/K is Galois if and only if the map $m : L \otimes_K L \to L^G$ is an isomorphism

Proof. (1) Assume that L/K is separable. We may choose an embedding of L in E and view L as a subfield. Then by Exercise 3.1.5, $L \otimes_K L$ is a subring of $E \otimes_K L$, so it suffices to show that $E \otimes_K L$ is reduced. If a is a nilpotent element of $A = E \otimes_K L$, any ring homomorphism from A to a field takes a to 0, so if there is a non-zero such a, homomorphisms cannot form a basis of $A^\vee$.

In the other direction, if $L \otimes_K L$ is reduced, then, by the remark above, it is a finite product of finite extensions of L. Hence $E \otimes_K L$ is a finite product of copies of E (since E is algebraically closed).

(2) Assume that m is an isomorphism. The product L^G is reduced, hence so is $L \otimes_K L$, so by the previous part, L/K is separable. The set of K-algebra maps from $L \otimes_K L$ to E is identified with the set of pairs of K-algebra maps from L to E (since $L \otimes_K L$ is the coproduct), so with $X \times X$. On the other side, $L^G = L \otimes_K K^G$, so K-maps from it to E coincide with pairs (x, g) with $x \in X$ and $g \in G$. Composition with m thus determines a map from $G \times X$ to $X \times X$, which is easily seen to be given by the projection and the action of G on X. Since m is an isomorphism, this map is a bijection. Thus L/K is also normal.

In the other direction, since L/K is separable, we have that $L \otimes_K L$ is a finite product of finite field extensions of L. Each such finite field extension can be embedded in E, so in particular determines an embedding of L in E. Since L is normal, we get that they all factor via an element of G. □

Let L be a Galois extension. We just saw that $L \otimes_K L$ is isomorphic to $L \otimes_K K^G$. In other words, after base change to L, the K-algebras L and K^G become isomorphic. Obviously, they are not isomorphic, as algebras, before the base change. However, the action of G on L makes it a representation of G.

Exercise 3.6.7. Show that if $L \otimes_K L$ is viewed as a base change of the representation L, then $m : L \otimes_K L \to L^G$ becomes an isomorphism of representations (where L^G is the regular representation). Show that m is not the base change of an isomorphism over K.

One may still ask if there is *some* isomorphism of L with the regular representation over K. We will see the answer later.

End of lecture 11, Apr 27

4. The Jacobson radical

In §2 we achieved a reasonably good understanding of semisimple rings. However, we had also seen that the class of such rings is rather small, and beyond group algebras and endomorphism algebras we saw no examples. Our current goal now is to pass to a larger class of rings, while maintaining similar properties. This class is defined in terms of the *Jacobson radical*, which makes the definition more elementary as well.

4.1. **Definition and basic properties.** Let M be a module over a ring R. If M is semisimple, it is a direct sum of simple modules, and in particular, every non-zero element of M has a non-zero projection to some simple module. In the language of the following definition, the radical of M is 0:

Jacobson radical
$J(M)$
Jacobson semisimple
JSS

Definition 4.1.1. The *Jacobson radical* $J(M)$ of a module M over a ring R is the collection of all elements $m \in M$ such that $f(m) = 0$ for all maps $f : M \to N$ where N is a simple module. M is called *Jacobson semisimple* (*JSS*) if $J(M) = 0$.

The Jacobson radical of R is its Jacobson radical as a module over itself.

Proposition 2.3.8 immediately provides the following more concrete description for rings:

Corollary 4.1.2. *The Jacobson radical of a ring is the intersection of all its maximal left ideals*

The discussion preceding the definition shows that a semisimple module is Jacobson semisimple. We will usually say "radical" in place of "Jacobson radical", but we cannot say "semisimple" in place of "Jacobson semisimple":

Example 4.1.3. The ring $\mathbb{Z}$ is Jacobson semisimple: any non-zero element is not divisible by some prime. We saw in Example 2.3.4 that it is not semisimple □

Example 4.1.4. If R is a commutative algebra finitely generated over a field k, then its Jacobson radical is the set of nilpotent elements in R. Indeed, the collection of simple modules consists of fields extensions of k which are finite (by Hilbert's Nullstellensatz). Any nilpotent element goes to 0 in any field, while if $a \in R$ is not nilpotent, the localisation $R[\frac{1}{a}]$ is again a non-zero finitely generated k-algebra, and each of its simple modules is also a simple module for R that does not contain a.

Hence, such a commutative algebra is Jacobson semisimple if and only if it is reduced. On the other hand, we had seen in Example 2.3.5 that $k[x]$ is not semisimple (and by Corollary 2.4.20, the same is true whenever R is infinite dimensional as a vector space over k). □

End of lecture 12, May 3 (There was some discussion re equivalence of categories, that I neglected before Exercise 2.4.8)

Example 4.1.5. Still in the commutative world, a local ring is Jacobson semisimple if and only if it is a field (recall that a local ring is a ring with a unique maximal ideal, which is therefore its Jacobson radical. For example, $k[[x]]$, the power-series ring) □

Example 4.1.6. If k is an algebraically closed field of characteristic p, and G is a p-group, then the Jacobson radical of the group algebra $R = k[G]$ is the *augmentation ideal*, the kernel of the map $\epsilon : R \to k$ sending each $g \in G$ to $1 \in k$. Indeed, a simple module over R is the same as an irreducible representation of G, and we saw in Proposition 2.5.12 that such a representation must be trivial, i.e., R acts via ϵ. □

augmentation ideal

We now proceed with establishing some basic properties of Jacobson radicals, parallel to §2.3 (but with less success):

Proposition 4.1.7. *Let M be a module over a ring R*

1. *If $f : M \to N$ is a homomorphism, $f(J(M)) \subseteq J(N)$ (in other words, J is a functor). In particular, a submodule of a Jacobson semisimple module is itself semisimple.*
2. *If M_α is a family of modules, $J(\oplus_\alpha M_\alpha) = \oplus_\alpha J(M_\alpha)$. Hence, the direct sum of Jacobson semisimple modules is JSS.*
3. *$M/J(M)$ is JSS, and is universal among maps from M to a JSS module.*
4. *$J(R)$ is a two-sided ideal, and $R/J(R)$ is a JSS ring.*
5. *$J(R)M \subseteq J(M)$*

Proof.

1. Let $a \in J(M)$ and let $t : N \to S$ be a map, where S is simple. Then $t \circ f : M \to S$ is a map from M to a simple module, hence $t(f(a)) = 0$. Thus $f(a) \in J(N)$.
2. We have $\oplus J(M_\alpha) \subseteq J(\oplus M_\alpha)$ by the previous part. On the other hand, again by the previous part we have $J(\prod M_\alpha) \subseteq \prod J(M_\alpha)$, and this clearly restricts to direct sums.
3. $J(M)$ is precisely the kernel of the product all maps from M to simple modules.
4. We already mentioned that $J(R)$ is an intersection of left ideals, so is a left ideal. On the other hand, it is clearly preserved by all endomorphisms, so it is a two-sided ideal by Corollary 2.4.4.

   By definition, R and $R/J(R)$ have the same simple modules. Hence, if $a \in R/J(R)$ maps to 0 in all simple modules over $R/J(R)$, its pre-image maps to 0 in all simple modules over R, so it belongs to $J(R)$.

(5) It suffices to show that $J(R)m \subseteq J(M)$ for all $m \in M$. This follows from (1), since $J(R)m$ is the image of the map $a \mapsto am$. □

Remark 4.1.8. Let us warn against some easily held misconceptions:

(1) In contrast with the semisimple case, a quotient of a JSS module need not be JSS. For example, let $R = k[x, y]/xy$ viewed as a module over itself, and let $N = R/x - y = k[x]/x^2$ (k a field). By Example 4.1.4, R is JSS, but $x \in N$ belongs to $J(N)$.

(2) Again in contrast with the semisimple case, R JSS does not imply that all R-modules are JSS (same example).

(3) If $f : R \to T$ is a map of rings, we may view T as an R-module, and so we have two Jacobson radicals for T, one as a T-module and one as an R-module. These need not be the same, does $J(R)$ map into $J(T)$. For example, if R is a non-trivial (commutative) local ring (say, $R = k[[x]]$), and T its fraction field, then the unique simple module is the residue field, and the unique map from T to it is 0, so $J_R(R)$ is the maximal ideal, $J_R(T) = T$ and $J_T(T) = 0$.

(4) For the same reason, $R/J(R)$ is not a universal JSS ring over R.

□

Our next goal is to give a more concrete description of the radical. This goes through a useful result called *Nakayama's Lemma*:

Proposition 4.1.9 (Nakayama's Lemma). *Let R be a ring.*

(1) Let M be a finitely generated module, and let $I \subseteq M$ be a submodule. The following are equivalent
   (a) $I \subseteq J(M)$
   (b) If $N \subset M$ is a proper submodule, then so is $N + I$

(2) Let $I \subseteq R$ be an ideal. The following are also equivalent
   (a) $I \subseteq J(R)$
   (b) If M is a finitely generated R-module, and $N \subset M$ is a proper submodule, then so is $N + IM$.
   (c) Each element of $1 + I$ is invertible on both sides.

Proof. (1) (1a) $\implies$ (1b): Assume $N + I = M$. Since M is finitely generated, there is a finitely generated $I_0 \subseteq I$ with $N + I_0 = M$. Hence we may assume that I is finitely generated (it is this assumption that we actually need). By induction, we may assume that I is generated by one element a.

By Zorn's lemma, we may assume N to be maximal among submodules contradicting the claim. We claim that M/N is simple. Otherwise, there is a proper submodule $N \subset L \subset M$, but since $N + Ra = M$, we cannot have $a \in L$, so L contradicts maximality of N. This shows that M/N is simple, but the image of a in it is non-zero, contradicting that $a \in J(M)$.

(1b) $\implies$ (1a)**:** Assume that $a \in I$ does not belong to some proper maximal submodule N. Then $M = N + Ra \subseteq N + I$, contradicting the assumption (this direction does not use that M is finitely generated).

(2) (2a) $\implies$ (2b)**:** Apply the previous part to the submodule $IM \subseteq J(M)$

(2b) $\implies$ (2c)**:** Fix $a \in I$, and let $M = R$ and $N = R(1 + a)$. Then $R = N + Ra \subseteq N + I$, hence $N = R$. It follows that $b(1+a) = 1$ for some $b \in R$, so $1 + a$ has a left inverse. On the other hand, $b = 1 - ba$ and $ba \in I$, so applying the same statement to ba in place of a, we find that b has both a right and a left inverse, which must agree (exercise) and thus b is also the right inverse of $1 + a$.

(2c) $\implies$ (2a)**:** If N is a left ideal in R and $N+I = R$, then $1 = n+a$ for $n \in N$ and $a \in I$. Then $n = 1 - a$ is invertible, so $N = R$. Thus, by the previous part, $I \subseteq J(R)$. □

End of lecture 13, May 4

It is often useful to view the direction (2a) $\implies$ (2b) above as stating a relation between a module over R and its reduction modulo I. Geometrically (in the commutative case), R can be thought of as the ring of functions on a neighbourhood X of the space X_0 whose ring of functions is R/I, and the result talks about extending from X_0 to X. For example:

Corollary 4.1.10. *If M is finitely generated over R, and the images of $m_1, \dots, m_k \in M$ in $\bar{M} = M/J(R)M$ generate $\bar{M}$, then $m_1, \dots, m_k$ generate M. In particular, if $\bar{M} = 0$ then $M = 0$.*

Proof. Let N be the submodule generated by the m_i, and apply Nakayama's Lemma. □

Example 4.1.11. The assumption that M is finitely generated is essential. For example, if R is a non-trivial commutative local ring and M is its fraction field, then $\bar{M} = 0$. □

Similarly, the third condition allows to related invertibility:

Corollary 4.1.12. *An element $x \in R$ is invertible if and only if $\bar{x} \in \bar{R} = R/J(R)$ is invertible*

Proof. Assume that $\bar{y}\bar{x} = 1$ for some $y \in R$. Then $yx = 1 + a$ for some $a \in J(R)$. By Nakayama's Lemma, $1 + a$ is invertible, so $(1 + a)^{-1}y$ is an inverse of x. The proof for the right inverse is the same, and the other direction is always true. □

The third condition can also be used to compute the radical: $x \in J(R)$ if and only if $1+ax$ is (two-sided) invertible. It also has the following surprising corollary (compare with similarly surprising Corollary 2.4.21, which was proved in a completely different way):

Corollary 4.1.13. *For any ring R,*

$$J(R) = \{x \in R \,|\, 1 + ax \in R^* \,\forall a \in R\} = \{x \in R \,|\, 1 + xa \in R^* \,\forall a \in R\}$$

In particular, $J(R) = J(R^\circ)$, and R is JSS if and only if R° is.

Proof. Assume that $x \in J(R)$. Since $J(R)$ is a two-sided ideal, $ax, xa \in J(R)$ for all $a \in R$. Hence, by Nakayama's Lemma, $1 + xa$ and $1 + ax$ are invertible for all $a \in R$. Conversely, if $1 + ax \in R^*$ for all a, then $I = Rx$ is contained in $J(R)$, so in particular $x \in R$. By the first direction, it follows that $1 + xR \subseteq R^*$. Applying this argument in R°, we find that the condition is symmetric. □

Remark 4.1.14. The last corollary also implies that the Jacobson radical is *definable* in the language of rings (in the sense of first-order logic), and therefore that the class of JSS rings is first-order axiomatizable (in contrast with semisimple rings). □

4.2. **Structure theory.** We had now convinced ourselves that JSS rings are easier to detect and to find than semisimple rings (which are a special case). For semisimple rings, however, we had a detailed structure theory: Each such ring is isomorphic to a finite product of semisimple primitive rings, and any primitive ring is isomorphic to the endomorphism ring of a finite dimensional vector space over a division ring. We now ask how much of this structure is retained.

Let S be a simple R-module. If S is not faithful, there are some elements $r \in R$ such that $rS = 0$. The collection of all such elements forms a two-sided ideal called the *annihilator* of S, $\operatorname{ann}(S)$. Then S is a faithful module over $R/\operatorname{ann}(S)$, which is still simple, so $R_S := R/\operatorname{ann}(S)$ is a primitive ring, with faithful simple module S. The product of these maps gives a map from R to a product of primitive rings. It turns out that the kernel is $J(R)$:

annihilator

Proposition 4.2.1. *For any ring R, the Jacobson radical $J(R)$ is the intersection of* $\operatorname{ann}(S)$ *over all simple R-modules S. In particular, R is JSS if and only if it embeds to a product of primitive rings with surjective projections.*

The example of $R = \mathbb{Z}$ shows that we cannot do better.

Proof. We already saw that $J(R)$ is the intersection of all kernels of R-modules maps $R \to S$, with S a simple R-module, so to prove the statement, it is enough to show that $\operatorname{ann}(S)$ is the intersection of all kernels of R-module maps to S. If $rS = 0$ and $f : R \to S$ is a map, then $f(r) = f(r1) = rf(1) = 0$. Conversely, for any $s \in S$ there is a map $f : R \to S$ with $f(1) = s$, so $rs = 0$. □

Next, we would like to understand primitive rings. We had already noticed (in Example 2.4.13) that for every division ring D and every D-module V, the ring $\operatorname{End}_D(V)$ is primitive, with V a faithful simple module. We also saw that in the semisimple case, these are the only ones. What happens in the general case?

Example 4.2.2. Let D be a division ring, let V be a D-module, and let $A \subseteq \operatorname{End}_D(V)$ be a sub-ring. Then V is a faithful A-module, and it is simple if for any $0 \neq u, v \in V$ there is $a \in A$ such that $au = v$.

Assume now that D is a countable algebraically closed field, and that V is a space with a countable basis B over D. For every finite subset $B_0 \subset B$, let $E(B_0)$ be the endomorphism ring of the subspace spanned by B_0, viewed as a (non-unital) subring of $\operatorname{End}_D(V)$. Let E be the union of all $E(B_0)$, and let $R = D + E \subseteq \operatorname{End}_D(V)$ (where D is embedded as the centre). Note that the product of two elements of E, or of an element of E with an element in D is again in E. Thus R is a subring. It is easily seen to act transitively on the non-zero elements of V, and so R is primitive.

We claim that R is not an endomorphism algebra. Indeed, any endomorphism of V over R must commute with all elements of $\operatorname{End}_D(V)$, so must be in D. Hence, we would have $R = \operatorname{End}_D(U)$ for some U. But R is countable, so U must be finite dimensional, and that is clearly impossible. □

The ring R in the last example was not the whole $\operatorname{End}_D(V)$, but in some sense it was "large". To explain this notion of largeness, we will make a small digression on function space topology.

End of lecture 14, May 10

4.2.3. *Topology on function spaces.* For spaces X and Y, we would like to produce a new space Y^X which behaves like the space of functions from X to Y, in the following sense: A (continuous) map u from a $Z \times X$ to Y (where Z is a third space) can be thought of as a family of maps from X to Y, parametrised by Z, and so should correspond to a (continuous) map $u^\vee : Z \to Y^X$. This should happen for every space Z uniformly: If $t : Z_1 \to Z$ is a map, the composed map $u \circ (t \times 1) : Z_1 \times X \to Y$ should correspond to the composed map $u^\vee \circ t$. This property determines Y^X uniquely: it is the final object in a suitable category (this is an exercise, but will be elaborated below).

If $h : X_1 \to X$ and $g : Y \to Y_1$ are maps of spaces, then we have induced maps $h \times 1 : X_1 \times Y^X \to X \times Y^X \xrightarrow{\text{ev}} Y$ corresponding to a map $(\circ h) : Y^X \to Y^{X_1}$ and similarly $(g\circ) : Y^X \to {Y_1}^X$. In other words, $(X, Y) \mapsto Y^X$ is a functor of Y, and a contravariant functor of X.

Applying the defining property in the case $Z = *$ (one point space), we find that the points (maps from $*$ to Y^X) are continuous maps from $X = * \times X$ to Y. This describes Y^X as a set, so to obtain it as a space, we need to define a topology. Further, applying the definition to the identity map on Y^X, we obtain a map $\text{ev} : Y^X \times X \to Y$, with the property that the map $u^\vee \mapsto u$ is determined by: $u = \text{ev} \circ (u^\vee \times 1)$. In particular, if $Z = *$, we have for each $x \in X$, $u(x) = \text{ev}(u, x)$, so ev is the evaluation map. Likewise, the identity map $X \times Y \to X \times Y$ corresponds to a continuous map $\text{ce} : X \to (X \times Y)^Y$, which is easily seen to be given by $x \mapsto (y \mapsto (x, y))$. Again we have the relation $u^\vee = (u\circ) \circ \text{ce}$, where $(u\circ)$ is as in the previous paragraph. The fact that these maps are inverse comes from the fact that the compositions $\text{ev} \circ (1 \times \text{ce})$ and $(\text{ev}\circ) \circ \text{ce}$ are the identity.

Thus, our goal is to determine a topology on the set Y^X of continuous maps from X to Y, such that $\mathrm{ev} : Y^X \times X \to Y$ and $\mathrm{ce} : X \to (X \times Y)^Y$ and the composition maps are continuous. To see what this topology might be, consider a point $x \in X$. The resulting composition $Y^X \to Y^* = Y$ should be continuous, so if $U \subseteq Y$ is open, the inverse image $\{f \in Y^X \mid f(x) \in U\}$ should be open as well.

More generally, still assuming that the topology on Y^X was defined, let $K \subseteq X$ be compact, and let $U \subseteq Y$ be open. If $f : X \to Y$ is continuous, and $f(K) \subseteq U$, then $K \times \{f\}$ is contained in the open set $V = \mathrm{ev}^{-1}(U)$. Hence, for every $x \in K$ there is an open neighbourhood $V_x \times W_x$ of $(x, f) \in K \times Y^X$ inside V. Since K is compact, it is contained in a finite union $V_{x_1} \cup \cdots \cup V_{x_m}$. $W = W_{x_1} \cap \cdots \cap W_{x_m}$ is an open neighbourhood of f in Y^X such that for all $g \in W$, $g(K) \subseteq U$.

In other words, for every compact $K \subseteq X$ and open $U \subseteq Y$, the subset $W(K, U) = \{f \in Y^X \mid f(K) \subseteq U\}$ is open. Note that if $g : Y \to Y_1$ and $h : X_1 \to X$ are continuous, $U_1 \subseteq Y_1$ is open, and $K_1 \subseteq X_1$ is compact, then $(g\circ)^{-1}(W(K, U_1)) = W(K, g^{-1}(U_1))$, and $(\circ h)^{-1}(W(K_1, U)) = W(h(K_1), U)$. Hence, we may try to put on Y^X the topology generated by such sets. This is called the *compact-open topology* (or the *uniform convergence topology*), and it does work for large classes of spaces (for example, the *compactly generated* ones, after some modifications).

compact-open topology

uniform convergence topology

Let us spell out what is the meaning of density in the compact-open topology: A set B of continuous functions from X to Y is dense if for any continuous function $f : X \to Y$, compact $K \subseteq X$ and open $U \subseteq Y$ containing $f(K)$, there is $g \in B$ with $g(K) \subseteq U$. In particular, for discrete Y, we should have, for each compact $K \subseteq X$, a function $g \in B$ that agrees with f on K. And if X is discrete as well, then "compact" means "finite".

4.2.4. *The density theorem.* We are now ready to state (and prove) Jacobson's density theorem. It is actually more convenient to prove a more general statement.

Theorem 4.2.5. *Let V be a faithful semisimple module over a ring R, and let $D = \mathrm{End}_R(V)$. Then R is a dense subring of $\mathrm{End}_D(V)$.*

Proof. We need to show that if $v_1, \dots, v_k \in V$ and T is an endomorphism of V over D, then there is $r \in R$ with $rv_i = T(v_i)$ for all i.

We first deal with the case $k = 1$. By semisimplicity, we may write $V = Rv_1 \oplus M$ for some submodule M. Since this is a direct sum as R-modules, the projection maps are R-module maps, i.e., they are in D. Hence T respects this decomposition, and in particular $T(Rv_1) \subseteq Rv_1$. Hence, $Tv_1 = rv_1$ for some $r \in R$.

For the general case, apply the case $k = 1$ to the semisimple module V^k and the endomorphism T in each coordinate. □

Corollary 4.2.6 (Jacobson's density theorem)**.** *Any primitive ring is isomorphic to a dense subring of* $\operatorname{End}_D(V)$ *for a division ring* D *and a vector space* V *over it.*

More explicitly, for any linearly independent $v_1, \dots, v_k \in V$ and any $u_i \in V$ there is $r \in R$ with $rv_i = u_i$. In other words, R acts transitively on finite independent tuples.

Proof. This is just the special case of the theorem, with V a simple faithful module over R. D is a division ring in this case, by Schur's Lemma □

Corollary 4.2.7. *Assume that* R *is primitive. Then either it is isomorphic to* $M_n(D)$ *for some division ring* D*, or for each* n *there is a subring* $R_n \subseteq R$ *and a surjective map from* R_n *to* $M_n(D)$*.*

Of course, R is semisimple if and only if the first case occurs.

Proof. Let V and D be as above. For subspace $U \subseteq V$, let $R_U = \{r \in R \mid rU \subseteq U\}$. Then R_U is a subring, and restriction defines a ring homomorphism $R \to R_U$. If U is finite dimensional, the density theorem assures that this map is surjective. □

Remark 4.2.8. This corollary recovers Theorem 2.4.16, this time with the natural action on V (cf. Remark 2.4.18). □

Corollary 4.2.9. *Any JSS ring is isomorphic to a subring of a direct product of endomorphism rings, with dense projections.*

End of lecture 15, May 11

As an example application, we prove the following (slightly contrived) result (from [Her68], where more interesting results of this kind can be found):

Proposition 4.2.10. *If* R *is a ring satisfying* $x^3 = x$ *for all commutators* x*, then* R *is commutative.*

Proof. We proceed by increasingly more general classes of rings. First, assume that R is a division ring. If $x = ab - ba$ is a non-zero commutator, we have $x^3 = x$, hence $x^2 = 1$, so x is 1 or -1, and in particular is central. Applying the equation to the commutator of a and ab, we find that $a(ab) - (ab)a = ax$ is in the centre as well, and since x is invertible, so is a. Hence, $x = 0$.

Next, we note that if V is a vector space of dimension 2 over a division ring D, then $\operatorname{End}_D(V)$ does not satisfy the assumption: we had seen (in Lemma 2.5.7) that an order 2 element in such a ring is a commutator.

Now assume that R is JSS. Then by Corollary 4.2.9, R embeds in some product $\prod R_\alpha$, where each R_α is an endomorphism ring, and R has dense image. If some R_α is not commutative, then R has a subring R_0 whose image is the endomorphism ring of a two-dimensional space. But the assumptions of the statement are preserved under both subrings and quotient, so we contradict the previous case.

Finally, let R be any ring. Then, the image of any x as in the assumption is 0 in $R/J(R)$, so $x \in J(R)$. But then both $1+x$ and $1-x$ are invertible, and we have $0 = x^3 - x = x(x-1)(x+1)$, so $x = 0$. □

4.3. **Chain conditions.** The next question we would like to answer is: which JSS rings are actually semisimple? The preceding examples suggest that the semisimple rings should be the JSS rings that are finite-dimensional in some sense. This turns out to be true, but the precise meaning of "finite-dimensional" should be clarified. As usual, it is most fruitful to work with general modules.

Artinian **Definition 4.3.1.** A module M is *Artinian* if there is no infinite descending chain of submodules $M = M_0 \supset M_1 \supset \dots$.

Noetherian A module M is *Noetherian* if there is no infinite ascending chain of submodules $0 = M_0 \subset M_1 \subset \dots \subset M$.

finite length M has *finite length* if the length of any chain of proper inclusions $0 \subset M_1 \subset \dots \subset M$ is bounded by a natural number (the least such number is
length of M then the *length of M*)

The ring R is Artinian or Noetherian or finite length, if it is such as a module over itself.

Example 4.3.2. If k is a field, the polynomial algebra $R = k[x]$ is Noetherian: any submodule (i.e., ideal) can be generated by a polynomial, and $Rp(x) \subseteq Rq(x)$ precisely if q divides p. Thus, the length of an increasing chain of ideals starting with $Rp(x)$ is bounded by the degree of p. It is not Artinian: the ideals generated by x^n, as n ranges over natural numbers, form an infinite descending sequence. Similarly, $\mathbb{Z}$ is Noetherian but not Artinian.

Geometrically, an ascending chain of ideals in an algebra of functions $k[X]$ corresponds to a descending chain of Zariski closed subsets of X, so the Noetherian condition corresponds to finite dimensionality. For example, $k[x]$ is the ring of functions on the affine line, and a proper ideal $Rp(x)$ corresponds to a finite set of points (with multiplicities), so a descending sequence must stabilise. The Artinian condition corresponds to ascending chains of closed subsets, so to satisfy it, X should be 0-dimensional (and then, the length of a chain is roughly the number of points). □

Example 4.3.3. To obtain a module that is Artinian but not Noetherian, note that if R is a k-algebra, and M is a module over it, then M is a vector space over k, and the k-dual $M^\vee$ is a module over R°, with $(r\varphi)(m) = \varphi(rm)$. If $N \subseteq M$ is a submodule, then $N^\perp = \{\varphi \in M^\vee \mid \varphi(N) = 0\} \subseteq M^\vee$ is also a submodule, and if $N_1 \subset N_2$ then ${N_2}^\perp \subset {N_1}^\perp$. In particular, if M is not Artinian, then $M^\vee$ is not Noetherian.

Taking $M = R = k[x]$, we find a non-Noetherian module $M^\vee$. This module is "too large" to be Artinian: for example, if $A \subseteq k$ is any subset, the subspace

$$N_A = \{\varphi \in M^\vee \,|\, (\prod_i^m (x - a_i))\varphi = 0 \text{ some } a_i \in A\}$$

of elements whose support lies in A is a submodule, and $N_A \subset N_B$ if $A \subset B$. However, if we restrict to the submodule $N_0 = \{\varphi \in M^\vee \,|\, x^m\varphi = 0\}$ of elements supported at 0, the same argument shows that N_0 is not Noetherian, and it is now Artinian: the submodules are all of the form $(k[x]/x^m)^\vee$ for some m. □

Proposition 4.3.4. *Let R be a ring*

(1) *If $N \subseteq M$ are modules, then M is Artinian or Noetherian if and only if N and M/N are.*
(2) *If R is a k-algebra, where k is a field, and M is an R-module that is finite-dimensional over k, then R has finite length.*
(3) *M is Noetherian (Artinian) if and only if every non-empty collection of submodules of M has a maximal (minimal) element.*
(4) *A module M has finite length if and only if it is both Noetherian and Artinian.*
(5) *M is Noetherian if and only if every submodule is finitely generated.*
(6) *R is Noetherian or Artinian or has finite length if and only if each finitely generated module is the same*

Proof. (1) An infinite chain of either kind in N is also such in M, and an infinite chain in M/N can be lifted to the same in M.

In the other direction, given an infinite chain M_α in M (of either kind), consider the sequence $M_\alpha \cap N$ in N. If it stabilises, we may assume that all $M_\alpha \cap N$ are equal to one module N_0 that is contained in all M_α. Then $M_\alpha/N = M_\alpha/N_0$ is a non-trivial sequence.

(2) Any infinite chain is also an infinite chain of k-vector spaces
(3) This is a statement about posets that follows directly from Zorn's Lemma.
(4) Assume M is Artinian, and let $M_0 = 0$. For each $i \in \mathbb{N}$, if $M_i \neq M$, let M_{i+1} be a minimal element among modules strictly containing M. If M is also Noetherian, this process has to stop at a finite stage. Now, if N_i is any chain of strict inclusions, the map that sends N_i to the least M_j containing it is injective (by construction of M_i), so its length is at most the number of M_i. The other direction is trivial.
(5) Assume M is Noetherian. Define $M_0 = 0$ and M_{i+1} the module generated by M_i and some element m_i not in M_i. This chain must stop, so the m_i generate M. Since a submodule of a Noetherian module is Noetherian, this direction is done.

In the other direction, let M_i be an ascending chain, and let N be its union. By assumption, it is generated by a finite number of elements, hence is equal to a finite subunion. Therefore, the chain stabilises.

(6) If M is finitely generated, it a quotient of a finite direct sum of copies of R, so the result follows from the previous parts. The other direction is trivial. □

Corollary 4.3.5. *A ring is semisimple if and only if it is Artinian and JSS*

Proof. If R is semisimple, it is a finite sum of simple modules, hence has finite length, so is Artinian, and we already saw it is JSS.

Assume that R is Artinian and JSS. Then finitely many maximal ideals have 0 intersection, so R embeds in a finite direct sum of simple modules. □

Thus, the difference between JSS and semisimple is the Artinian condition. We now drop the JSS assumption and ask: What can be deduced about the Jacobson radical if we assume that the ring is Artinian?

nilpotent ideal

Recall that an element a of R is nilpotent if $a^n = 0$ for some n. An ideal $I \subseteq R$ is a *nilpotent ideal* if $I^n = 0$ for some n. It is clear that each element of a nilpotent ideal is nilpotent, but the converse may fail, for two different reasons: First, it might happen that the n required for different elements grows uboundedly, and second, the product of two nilpotent elements need not be nilpotent. Indeed, even an ideal generated by one nilpotent element may contain elements that are not nilpotent (the second problem does not occur if R is commutative, and then the first does not occur either if R is Noetherian, so for such rings, an ideal is nilpotent if and only if all its elements are).

nil ideal

An ideal consisting entirely of nilpotent elements is called a *nil ideal*. It is easy to check that such an ideal I is contained in the radical: if $x \in I$, the standard formal power series representing $\frac{1}{1+x}$ is, in fact, a polynomial in x. For Artinian rings, the converse of all these statements is true:

Proposition 4.3.6. *If R is Artinian, then $J(R)$ is nilpotent, and an ideal is nilpotent if and only if it is nil, if and only if it is contained in $J(R)$.*

Proof. Let $J = J(R)$. The descending sequence J^k stabilises, so assume that $J^k = J^{k+1}$. In particular, the collection of non-zero ideals I such that $I = JI$ is non-empty. Let I be a minimal such ideal. Since $I \neq 0$, there is some $0 \neq x \in J^k I = I$. Then $J^k x$ is a subideal of I satisfying the defining condition, and thus I is generated by x. In particular, I is finitely generated, so by Nakayama's Lemma, $JI = I$ implies $I = 0$.

For the second statement, we mentioned above that if I is nil then $I \subseteq J$, and if $I \subseteq J$ then it is nilpotent by the first part. □

We may now return a debt:

Proof of Prop. 3.6.5. A is finite dimensional, hence Artinian, so semisimple is equivalent to JSS, and we just saw that $J(A)$ is the set of nilpotents. □

The proposition suggests the following definition:

semiprimary

Definition 4.3.7. A ring R is *semiprimary* if $J(R)$ is nilpotent, and $\bar{R} = R/J(R)$ is Artinian.

The interest in such rings comes from the following result:

Proposition 4.3.8. *Suppose that R is semiprimary, and that C is a class of modules over R that contains all simple modules, and is closed under exact sequences. Then C contains all Artinian modules.*

"Closed under exact sequences" means that for $N \subseteq M$, M is in C if and only if N and M/N are.

Proof. Let J be the radical of R. By assumption, $J^k = 0$ for some k, so for every module M over R, there is a minimal m such that $J^m M = 0$. Do induction on m. For the base, assume that $JM = 0$. Then M is actually a module over R/J which, by assumption, is Artinian, and thus semisimple. Since M is Artinian, it is a finite direct sum of simple modules, so is in C.

For $m > 1$, $JM \subseteq M$ and M/JM are Artinian with lower value of m, so are in C by induction. Hence, so is M □

We hadn't seen examples of Artinian rings which are not Noetherian. The reason is:

Corollary 4.3.9 (Akizuki–Levitzki–Hopkins)**.** *Any Artinian module over a semiprimary (in particular, Artinian) ring is Noetherian. In particular, an Artinian ring is Noetherian.*

Proof. Apply Proposition 4.3.8 to the class of Noetherian modules (which satisfies the assumptions by Proposition 4.3.4). □

End of lecture 16, May 17

5. Categories of modules

A general reference for this topic is [Bas68, Ch. 1, 2]

5.1. **Functors between module categories.** Let $t : R \to S$ be a map of rings. This map determines, as we saw in §3.4, a functor $t_* : R - \mathcal{M}od \to S - \mathcal{M}od$, given by tensoring with S over R (i.e., base change to S). Of course, this is not the only possible functor between these categories, and we may ask: Given an abstract functor $\mathbf{F} : R - \mathcal{M}od \to S - \mathcal{M}od$, can we recognise, from categorical considerations, that it is, in fact, t_*?

More generally, if N is a module over R°, we have a functor $\mathbf{F}_N$ on $R - \mathcal{M}od$ given on objects by $M \mapsto N \otimes_R M$. In general, it lands in $\mathcal{A}b$, but if N also has a structure of an S-module, which commutes with the R-structure (i.e., $(sn)r = s(nr)$ for all $r \in R$, $s \in S$ and $n \in N$) then we obtain a functor from $R - \mathcal{M}od$ to $S - \mathcal{M}od$. Such an N is called an *$S - R$ bi-module*, and the structure is equivalent to the structure of a module over $S \otimes R^\circ$. Again we may ask:

$S - R$ **bi-module**

Question 5.1.1. Which functors $\mathbf{F} : R - \mathcal{M}od \to S - \mathcal{M}od$ are of the form $\mathbf{F}_N = N \otimes -$ for some module N over $S \otimes R^\circ$? □

5.1.2. *Categories of functors.* To answer this question, we first need to explain more precisely what do we mean by "of the form". Consider, for example, the case $R = \mathbb{Z}$, so that $R - \mathcal{M}od = \mathcal{A}b$. For each prime p, there is a functor that sends an abelian group A to the quotient A/pA. Since the resulting group is actually an $\mathbb{F}_p$ vector-space, we have a canonical map $\mathbb{F}_p \otimes A \to A/pA$, and it is easy to check that this is an isomorphism. The fact that the isomorphism is canonical means that if $f : A \to B$ is a homomorphism, the isomorphism is compatible with the induced maps $A/pA \to B/pB$ and $\mathbb{F}_p \otimes A \to \mathbb{F}_p \otimes B$. Thus, despite literally being different, the functor $A \mapsto A/pA$ is "essentially" base change to $\mathbb{F}_p$. This is captured more precisely by saying that the functors are isomorphic, as objects in the category of functors:

natural transformation

Definition 5.1.3. Let $\mathbf{F}, \mathbf{G} : \mathcal{C} \to \mathcal{D}$ be functors between categories $\mathcal{C}$ and $\mathcal{D}$. A *natural transformation* α from $\mathbf{F}$ to $\mathbf{G}$ consists of a collection of maps $\alpha_A : \mathbf{F}(A) \to \mathbf{G}(A)$ for each object A of $\mathcal{C}$, such that

$$\begin{tikzcd}
\mathbf{F}(A) \arrow[r, "\mathbf{F}(t)"] \arrow[d, "\alpha_A"] & \mathbf{F}(B) \arrow[d, "\alpha_B"] \\
\mathbf{G}(A) \arrow[r, "\mathbf{G}(t)"] & \mathbf{G}(B)
\end{tikzcd} \tag{5.1}$$

commutes for all morphisms $t : A \to B$ in $\mathcal{C}$.

category of functors

$\mathcal{F}unct(\mathcal{C}, \mathcal{D})$

The category whose objects are functors from $\mathcal{C}$ to $\mathcal{D}$, and whose morphisms are natural transformations is called the *category of functors* between $\mathcal{C}$ and $\mathcal{D}$, denoted $\mathcal{F}unct(\mathcal{C}, \mathcal{D})$.

We had already seen some examples of natural transformations:

Exercise 5.1.4. Let $\mathbf{F}$ and $\mathbf{G}$ be the functors from $\mathcal{S}et$ to $\mathcal{G}p$ constructing the free and free abelian groups out of a set, respectively. Find an interesting natural map from $\mathbf{F}$ to $\mathbf{G}$.

Exercise 5.1.5. Find an interesting natural map from the identity functor on the category $k - \mathcal{M}od$ of k-vector spaces (k a field), to the functor $X \mapsto X^{\vee\vee}$. Show that when restricted to finite-dimensional spaces, this map is an isomorphism of functors.

Exercise 5.1.6. Show that a natural map α is an isomorphism if and only if each component α_A is an isomorphism

Exercise 5.1.7. Recall that if we view a monoid G as a category with one object, then a functor into $\mathcal{S}et$ is simply a G-set. Show that a map between such functors is, likewise, the same as a map of G-sets.

equivalence of categories

At this point, it is also convenient to define the notion of an *equivalence of categories.* As in any situation where the notion of a morphism is defined, we can define an isomorphism of categories to be a functor which is invertible.

Though this makes sense, this is not very useful, since it is too rigid. For example, we have seen in Exercise 2.4.8 that the category of modules over $R \times S$ is *essentially* the same as the category of pairs (M, N) of an R-module M and an S-module N. However, the categories are not actually isomorphic. Instead, the composition of the functors in both directions are *isomorphic* to the identity. This leads to the following definition:

Definition 5.1.8. Let $\mathcal{C}$ and $\mathcal{D}$ be two categories. An *equivalence* between $\mathcal{C}$ and $\mathcal{D}$ is a pair of functors $\mathbf{F} : \mathcal{C} \to \mathcal{D}$ and $\mathbf{G} : \mathcal{D} \to \mathcal{C}$, along with natural isomorphisms from $\mathbf{G} \circ \mathbf{F}$ and from $\mathbf{F} \circ \mathbf{G}$ to the appropriate identity functors.

$\mathcal{C}$ and $\mathcal{D}$ are *equivalent categories* if there is an equivalence between them **equivalent categories**

Often one refers just to either of the functors $\mathbf{F}$ or $\mathbf{G}$ as the equivalence.

Example 5.1.9. The construction in Exercise 2.4.8 determines an equivalence (the requirement that the maps α and β are natural is omitted in the exercise, but should hold for any reasonable solution of the exercise) □

Example 5.1.10. The assignment $V \mapsto V^\vee$ determines an equivalence from the category of finite dimensional vector spaces (over a fixed field) and its opposite. An inverse (up to isomorphism) is given by the same functor, viewed on the opposite categories. Note that again, the composition is not literally the identity (a vector space V is not *equal* to $V^{\vee\vee}$). □

Example 5.1.11. Exercise 5.1.7 can be summarised as saying that there is an equivalence between the category of G-sets, and the category of functors from G to $\mathcal{S}et$. □

Often, one direction of the equivalence is easier to define than the other. In such cases, it is useful to have the following fact:

Fact 5.1.12. *Let $\mathbf{F} : \mathcal{C} \to \mathcal{D}$ be a functor, and assume:*

(1) *The map $h \mapsto \mathbf{F}(h)$ from $\mathbf{Mor}_{\mathcal{C}}(x, y)$ to $\mathbf{Mor}_{\mathcal{D}}(\mathbf{F}(x), \mathbf{F}(y))$ is a bijection for all objects x, y of $\mathcal{C}$ ($\mathbf{F}$ is* fully faithful*)* **fully faithful**

(2) *For every object d of $\mathcal{D}$ there is an object c of $\mathcal{C}$ such that $\mathbf{F}(c)$ is isomorphic to d ($\mathbf{F}$ is* essentially surjective*)* **essentially surjective**

Then $\mathbf{F}$ is an equivalence of categories.

Example 5.1.13. Let k be a field, and let $\mathcal{C}$ be the category whose objects are the natural numbers, for $i, j \in \mathbb{N}$ the morphisms are $i \times j$ matrices over k, and composition is given by matrix multiplication. Then $\mathcal{C}$ is equivalent to the category $\mathcal{V}ec_k$ of finite dimensional vector spaces over k. The functor that attaches to $i \in \mathbb{N}$ the space k^i, and to each matrix the linear map it represents, satisfies the conditions of the fact. Note that the categories are not isomorphic: $\mathcal{C}$ has countably many objects, while in $\mathcal{V}ec_k$ the collection of objects is a proper class. □

Let R and S be fixed non-zero rings. We now have the language to state precisely our problem:

Question 5.1.14. Which functors $\mathbf{F} : R - \mathcal{M}od \to S - \mathcal{M}od$ are isomorphic to the functor $\mathbf{F}_N = N \otimes -$ for some module N over $S \otimes R^\circ$? □

We first note that not all functors are of this form:

Example 5.1.15. Let $\mathbf{F} : R - \mathcal{M}od \to S - \mathcal{M}od$ be the constant functor $M \mapsto S$ (all maps go to the identity). The there is no module N over $S \otimes R^\circ$ with $\mathbf{F}$ isomorphic to $\mathbf{F}_N$. Otherwise, we would have isomorphisms $0 = N \otimes_R 0 = \mathbf{F}_N(0) = \mathbf{F}(0) = S$. □

Example 5.1.16. Assume $R = S = \mathbb{Z}$, let $n > 1$ be a natural number, and let $\mathbf{F} : \mathcal{A}b \to \mathcal{A}b$ be the functor that assigns to A the kernel of multiplication by n. If $\mathbf{F} = \mathbf{F}_N$ for some abelian group N, then we have $0 = \mathbf{F}(\mathbb{Z}) = \mathbf{F}_N(\mathbb{Z}) = N \otimes \mathbb{Z} = N$, but $\mathbb{Z}/n\mathbb{Z} = \mathbf{F}(\mathbb{Z}/n\mathbb{Z}) = \mathbf{F}_N(\mathbb{Z}/n\mathbb{Z}) = N/nN$. □

The last example can be reformulated as follows: we have an abelian group A and a subgroup B, where $\mathbf{F}(A/B)$ is not a quotient of $\mathbf{F}(A)$, while we have:

Exercise 5.1.17. Show that if L is a submodule of M, then $N \otimes_R M/N \otimes_R L$ is canonically isomorphic to $N \otimes_R (M/L)$ (note that $N \otimes_R L$ need not be a submodule of $N \otimes_R M$, but there is a map)

The upshot of these examples is that there are certain categorical constructions that the functor $N \otimes -$ preserves, and should therefore be preserved by any functor isomorphic to it. We next study these constructions.

5.1.18. *Limits.* Suppose that A and B are subsets of a set X. Is there a categorical description of their intersection? The fact that A and B are subsets of X can be described by a "diagram" of morphisms

$$\begin{tikzcd} & B \arrow[d] \\ A \arrow[r] & X \end{tikzcd} \tag{5.2}$$

The intersection can be placed in the upper-left corner of the diagram: it has maps to both A and B, which agree when composed with the maps into X. Furthermore, it is universal with this property: any map to X that factors through both A and B, is in fact a map into $A \cap B$. Thus, $A \cap B$ is a universal completion of this diagram.

We would like to consider other universal completions of diagrams. Intuitively, such a diagram is given by some objects and arrows between them. Thus, a diagram in $\mathcal{C}$ is simply a functor from a fixed category $\mathcal{I}$ int $\mathcal{C}$. $\mathcal{I}$ is often called (and thought of as) the *index category*, and the diagram itself is said to be of type $\mathcal{I}$.

index category

For instance, in the example above, $\mathcal{I}$ is a category with 3 objects, a, b, x, and two morphisms, from a to x and from b to x (in addition to the identity morphisms). The diagram above is represented by the functor sending a to A, b to B and x to X. To express the universal property, we consider the category of "cones" over the diagrams:

Definition 5.1.19. Let $\mathbf{F} : \mathcal{I} \to \mathcal{C}$ be a diagram. The *slice category* for $\mathbf{F}$, $\mathcal{C}/\mathbf{F}$, is the category whose objects consist of an object X of $\mathcal{C}$, and for each object i of $\mathcal{I}$, a morphism $p_i : X \to \mathbf{F}(i)$, such that for each morphism $u : i \to j$ we have $\mathbf{F}(u) \circ p_i = p_j$. A morphism from (X, p) to (Y, q) in $\mathcal{C}/\mathbf{F}$ is a morphism $t : X \to Y$ in $\mathcal{C}$, such that $q_i \circ t = p_i$ for all objects i in $\mathcal{I}$.

slice category
$\mathcal{C}/\mathbf{F}$

The *limit* (or *inverse limit*) of the diagram $\mathbf{F}$, denoted $\varprojlim_{\mathcal{I}} \mathbf{F}$ is a final object in $\mathcal{C}/\mathbf{F}$.

The *colimit* (or *direct limit*) is the dual notion: it is the limit in $\mathcal{C}^\circ$ of $\mathbf{F}$, viewed as a functor from $\mathcal{I}^\circ$ to $\mathcal{C}^\circ$.

limit
inverse limit
$\varprojlim_{\mathcal{I}} \mathbf{F}$
colimit
direct limit

Of course, as with any universal object, the limit need not exist, but if it exists, it is unique. As before, we will often view the limit as an object of $\mathcal{C}$ (forgetting the morphisms p_i).

Example 5.1.20. The limit of the empty diagram is the terminal object □

Example 5.1.21. Let $\mathcal{I}$ be the set $\{0, 1\}$, viewed as a category (i.e., the objects are 0 and 1, and the only morphisms are identities). A diagram then consists of two objects, and the limit is their product (if exists). □

Example 5.1.22. Let $\mathcal{I}$ be the category with objects $0, 1$ and two morphisms from 0 to 1. A diagram then consists of two objects A, B and two morphisms f, g from A to B. The limit is then a map $h : K \to A$ universal with the property that $f \circ h = g \circ h$. It is called the *equaliser* of f and g. The colimit of the same diagram is called the *coequaliser* of f and g. □

equaliser
coequaliser

Example 5.1.23. Let $\mathcal{I}$ be the category with objects $0, 1, 2$, and morphisms $1 \to 0$ and $2 \to 0$. A diagram is then a configuration as in (5.2) (but the morphisms need not be inclusions). The limit of such a diagram is then an object in the upper left corner, which maps to A and to B, so that the two composed maps into X are the same, and which is universal among such. It is called the *fibred product* of A and B over X (though it of course depends on the morphisms, as well as on the objects), denoted $A \times_X B$ (it is also called the *pullback*). The dual notion of the colimit of two objects over another is called the *fibred coproduct* or *pushout*, denoted $A \coprod_X B$. □

fibred product
$A \times_X B$
pullback
fibred coproduct
pushout
$A \coprod_X B$

Exercise 5.1.24. Show that equalisers, coequalisers, fibred product and fibred coproducts all exist in the category $\mathcal{S}et$ of sets, and describe them explicitly.

When working with general categories, a colimit is simply a limit in the opposite category, so we mostly concentrate on limits below.

Exercise 5.1.25. Show that the following are equivalent for a category $\mathcal{C}$:

1. All (finite) products and all equalisers exist
2. There is a terminal object, and all fibred product exist

Exercise 5.1.26. Let G be a group, and X a G-set. Viewing X as a functor from G (i.e., a G-shaped diagram), What are the limit and colimit?

Exercise 5.1.27. Show that if $\mathcal{I}$ has an initial object i, then for any diagram $\mathbf{F} : \mathcal{I} \to \mathcal{C}$, the limit exists and is equal to $\mathbf{F}(i)$.

Exercise 5.1.28. There are at least two additional ways to describe the limit:

(1) Let $\mathcal{I}_*$ be the category $\mathcal{I}$ along with an additional object $*$, which has a unique morphism to each object. Any diagram of type $\mathcal{I}_*$ restricts to a diagram of type $\mathcal{I}$. Let $\mathbf{F}$ be a diagram of type $\mathcal{I}$. Then $\mathcal{C}/\mathbf{F}$ is equivalent to the category of diagrams of type $\mathcal{I}_*$ that restrict to $\mathbf{F}$, and morphisms that restrict to the identity.

(2) For any object X of $\mathcal{C}$, there is a constant diagram $\mathbf{F}_X : \mathcal{I} \to \mathcal{C}$, that attaches X to each object, and the identity to each morphism. Then $\mathcal{C}/\mathbf{F}$ is equivalent to the category of maps $\mathbf{F}_X \to \mathbf{F}$ and morphisms between them.

What limits exist in the category of modules?

Proposition 5.1.29. *For every ring R, the category $R-\mathcal{M}od$ has all limits and colimits.*

Proof. We first recall that we had already noted that the products and coproducts of an arbitrary set of modules exists, and is given by the direct product and direct sum, respectively.

Given $\mathbf{F} : \mathcal{I} \to R - \mathcal{M}od$, let

$$M = \{(m_i) \in \prod_{i \in \mathbf{Ob}(\mathcal{I})} \mathbf{F}(i) \,|\, \mathbf{F}(u)(m_i) = m_j \;\forall u \in \mathbf{Mor}_{\mathcal{I}}(i,j)\}$$

and let

$$N = \oplus_{i \in \mathbf{Ob}(\mathcal{I})} \mathbf{F}(i) / \langle \{m_j - \mathbf{F}(u)(m_i) \,|\, m_i \in \mathbf{F}(i), m_j \in \mathbf{F}(j), u \in \mathbf{Mor}_{\mathcal{I}}(i,j)\} \rangle$$

Then $M = \varprojlim \mathbf{F}$ and $N = \varinjlim \mathbf{F}$ (exercise) □

Remark 5.1.30. The expression $\langle \{m_j - \mathbf{F}(u)(m_i) \,|\, m_i \in \mathbf{F}(i), m_j \in \mathbf{F}(j), u \in \mathbf{Mor}_{\mathcal{I}}(i,j)\} \rangle$ that occurs in the proof can be viewed either as "the submodule generated by" or "the subgroup generated by". In other words, the subgroup is automatically a submodule (exercise) □

In the construction of the limit above, we really used just products and equalisers. This is not special to modules: one can use essentially the same construction to prove the following fact (a limit is finite if the index category has finitely many morphisms):

Proposition 5.1.31. *A category $\mathcal{C}$ admits all (finite) limits if and only if it admits all (finite) products and all equalisers.*

Exercise 5.1.32. Prove the proposition

5.1.33. *Preservation of limits.* Assume now that limits exist in categories $\mathcal{C}$ and $\mathcal{D}$. What happens when we apply a functor to a limit? First, if $\mathbf{G} : \mathcal{C} \to \mathcal{D}$ is a functor, and $\mathbf{F} : \mathcal{I} \to \mathcal{C}$ is a diagram, we obtain a diagram $\mathbf{G} \circ \mathbf{F}$ in $\mathcal{D}$. If X is the limit of $\mathbf{F}$ in $\mathcal{C}$, we have a structure map $p_i : X \to \mathbf{F}(i)$ for each $i \in \mathcal{I}$, and therefore maps $\mathbf{G}(p_i) : \mathbf{G}(X) \to \mathbf{G}(\mathbf{F}(i))$, so by the universal property, we have a morphism $\mathbf{G}(X) = \mathbf{G}(\varprojlim \mathbf{F}) \to \varprojlim \mathbf{G} \circ \mathbf{F}$.

Definition 5.1.34. Assume $\mathcal{C}$ and $\mathcal{D}$ have limits. A functor $\mathbf{G} : \mathcal{C} \to \mathcal{D}$ *preserves limits of type* $\mathcal{I}$ if for any diagram $\mathbf{F} : \mathcal{I} \to \mathcal{C}$, the canonical map $\mathbf{G}(\varprojlim \mathbf{F}) \to \varprojlim \mathbf{G} \circ \mathbf{F}$ is an isomorphism. It *preserves (finite) limits* if it preserves limits of type $\mathcal{I}$ for all (finite) $\mathcal{I}$.

preserves limits of type $\mathcal{I}$
preserves (finite) limits

Example 5.1.35. Let $f : R \to S$ be a map of rings. We have a functor $f^* : S - \mathcal{M}od \to R - \mathcal{M}od$, obtained by making an S module an R-module via f. This functor preserves all limits and colimits (by construction) □

Example 5.1.36. The functor that constructs a vector space from a set preserves all colimits, but not limits □

An argument similar to that of Proposition 5.1.31 shows:

Proposition 5.1.37. *A functor* $\mathbf{F} : \mathcal{C} \to \mathcal{D}$ *preserves all (finite) limits if and only if it preserves all (finite) products and all equalisers.*

Exercise 5.1.38. Prove it

Remark 5.1.39. In the category of modules, the set of morphisms between two objects is actually an abelian group. In particular, there is a 0-morphism between any two objects. The equaliser and coequaliser of a map f with the 0 map is, respectively, the kernel and cokernel (quotient map). Since composition with a fixed morphism is a group homomorphism, the (co)equaliser of two maps f and g is the same as the (co)kernel of $f - g$. Thus, arbitrary (co)limits can be constructed from (co)products and (co)kernels.

Likewise, if a functor respects the additive structure (i.e., the maps on Hom sets is a group homomorphism), then preservation of (co)equalisers is equivalent to preservation of (co)kernels. □

We now come to the result that motivated the discussion:

Proposition 5.1.40. *Let N be an $S \otimes R^{\circ}$-module for some rings R and S. Then the functor $R - \mathcal{M}od \to S - \mathcal{M}od$ given by $M \mapsto N \otimes_R M$ preserves all colimits.*

Before giving the proof, we make the following key observation:

Exercise 5.1.41. Let N, M and A be abelian groups, and let $\mathrm{Hom}(M, A)$ be the abelian group of homomorphisms from M to A. Given a homomorphism $t : N \to \mathrm{Hom}(M, A)$, the map $(n, m) \mapsto t(n)(m)$ is bilinear, and therefore determines a homomorphism $t' : N \otimes M \to A$.

(1) Show that $t \mapsto t'$ is an isomorphism

(2) Assume that M and N are right and left R-modules, respectively. Since $M \otimes N$ maps surjectively to $M \otimes_R N$, the homomorphisms $M \otimes_R N \to A$ can be viewed as a subgroup of $\operatorname{Hom}(M \otimes N, A)$. On the other hand, $\operatorname{Hom}(M, A)$ admits a (left) R-module structure, given by the right structure on M. Show that $\operatorname{Hom}(M \otimes_R N, A)$ corresponds to R-module homomorphisms $N \to \operatorname{Hom}(M, A)$.

(3) If M, in addition, has an S-module structure that commutes with the R°-module structure, the S-module maps from $M \otimes_R N$ to an S-module A correspond to R-module maps from N to $\operatorname{Hom}_S(M, A)$ (the subgroup of $\operatorname{Hom}(M, A)$ consisting of S-module homomorphisms). What happens if the module structures on M do not commute?

(4) In each case, both sides of the correspondence are functorial in each argument (when the other two are fixed). Show that the correspondence forms a natural isomorphism of functors

Proof of Prop. 5.1.40. By Example 5.1.35, we may ignore the S-structure, and view $N \otimes_R -$ as a functor into $\mathcal{A}b$.

Let $\mathbf{F}$ be a diagram of R-modules. Our task is to show that the map $t : \varinjlim N \otimes_R \mathbf{F} \to N \otimes_R \varinjlim \mathbf{F}$ is an isomorphism, i.e., we would like to construct an inverse s.

According to the exercise preceding the proof, to define the map s, we map equivalently define an R-module map $s' : \varinjlim \mathbf{F} \to \operatorname{Hom}(N, \varinjlim N \otimes_R \mathbf{F})$. But defining such a map is equivalent, by the definition of the colimit, to defining a collection of maps $\mathbf{F}(i) \to \operatorname{Hom}(N, \varinjlim N \otimes_R \mathbf{F})$ for each i (satisfying some compatibilities). This, in turn, is again equivalent to providing a collection of maps $N \otimes \mathbf{F}(i) \to \varinjlim N \otimes_R \mathbf{F}$, which we have by the definition of the colimit.

This constructs the map in the opposite direction. The verification that it is the inverse is left as an exercise. □

Remark 5.1.42. We note that the proof used very little information about modules. In fact, the statement is a special case of a purely categorical fact (with essentially the same proof) that we will discuss later. □

End of lecture 17, May 24

We note the last necessary condition covers the counterexamples we had so far. Thus we may imagine that this is also a sufficient condition. This is the contents of:

Theorem 5.1.43 (Eilenberg–Watts). *A functor $\mathbf{F} : R - \mathcal{M}od \to S - \mathcal{M}od$ is isomorphic to $\mathbf{F}_N$ for some $S \otimes R^\circ$-module N if and only if it commutes with all colimits (equivalently, with all quotients and direct sums).*

The “only if” part is Prop. 5.1.40, so let $\mathbf{F} : R - \mathcal{M}od \to S - \mathcal{M}od$ be a functor. To prove the theorem, we first need to answer the question: Who is N? To answer this, we note that $\mathbf{F}_N(R) = N \otimes_R R$ is canonically isomorphic to N as an $S \otimes R^\circ$-module, so it makes sense to set $N = \mathbf{F}(R)$.

Thus, N is an S-module, but what is the R°-structure? If φ is an endomorphism of R as an R-module, we get, by functoriality, an endomorphism

$\mathbf{F}(\varphi)$ of $N = \mathbf{F}(R)$, as an S-module. Thus, N is a module over $S \otimes \operatorname{End}_R(R)$. But we have seen in Prop. 2.4.2 that $\operatorname{End}_R(R) = R^\circ$.

So we wish to define an isomorphism of S-modules from $N \otimes_R M$ to $\mathbf{F}(M)$, and we already have it when $M = R$. We now note that every module M over R can be presented as $M = \oplus_\alpha R/f(\oplus_\beta R)$, i.e., the quotient of a free R-module by the image of a free module under some map f: First, choose some generators m_α for M and send 1 in the α-th copy of R to m_α. This provides a surjective map to M with a kernel $K \subseteq \oplus_\alpha R$. Now repeat the same for K in place of M to obtain the map f.

Given such a presentation $\oplus_\beta R \xrightarrow{f} \oplus_\alpha R \to M \to 0$, we may apply a functor $\mathbf{F}$ to obtain a sequence of maps $\mathbf{F}(\oplus_\beta R) \xrightarrow{\mathbf{F}(f)} \mathbf{F}(\oplus_\alpha R) \to \mathbf{F}(M) \to \mathbf{F}(0)$, and if $\mathbf{F}$ commutes with colimits, this sequence is still *exact*, i.e., the image of each map is the kernel of the following one, and we get a "presentation" $\mathbf{F}(M) = \oplus_\alpha \mathbf{F}(R)/\mathbf{F}(f)(\oplus_\beta \mathbf{F}(R))$. Hence, if we have two such functors, $\mathbf{F}_1$ and $\mathbf{F}_2$, a map $t : \mathbf{F}_1(R) \to \mathbf{F}_2(R)$ will induce a map will induce a map $t_M : \mathbf{F}_1(M) \to \mathbf{F}_2(M)$.

The problem with this approach is that the map appears to depend on the choice of presentation of M and so, while it works for an individual M, it is not clear that maps obtained in this way for different objects will amalgamate to a natural transformation. We will present two ways to overcome this difficulty, the first completely general (which we will only sketch), and the second uses the definition of tensor products.

Proposition 5.1.44. *Let $\mathbf{F}_1, \mathbf{F}_2 : R - \mathcal{M}od \to \mathcal{C}$ be two functors that commute with colimits (in particular, the apropriate colimits exist in $\mathcal{C}$). Then any map $t' : \mathbf{F}_1(R) \to \mathbf{F}_2(R)$ extends to a unique maps of functors $t : \mathbf{F}_1 \to \mathbf{F}_2$ (so that $t_R = t'$)*

Sketch of proof. For each R-module M, let $\mathcal{I}_M$ be the category whose objects are presentations $\oplus_\beta R \to \oplus_\alpha R \to M \to 0$, and whose morphisms are maps of diagrams which are the identity on M. Let $\pi_M : \mathcal{I}_M \to R - \mathcal{M}od$ be the functor that sends a presentation as above to the middle term $\oplus_\alpha R$. It is easy to see that $\varinjlim \pi_M = M$, and so we have $\mathbf{F}_i(M) = \varinjlim(\mathbf{F}_i \circ \pi_M)$. Now, a map $t_M : \mathbf{F}_1(M) \to \mathbf{F}_2(M)$ is determined uniquely by the restrictions to the items $\pi_M(X)$ in the diagram, and each such object is a direct sum of copies of R, so t_M is determined by the restrictions to the direct summands $\mathbf{F}_1(R)$, which is determined (and conversely, this restriction determines a map on $\mathbf{F}_1(M)$) □

The second proof constructs a map directly: Let $N = \mathbf{F}(R)$. By Exercise 5.1.41, constructing an S-module map $t'_M : N \otimes_R M \to \mathbf{F}(M)$ amounts to constructing an R-module map

$$t_M : M \to \operatorname{Hom}_S(N, \mathbf{F}(M)) = \operatorname{Hom}_S(\mathbf{F}(R), \mathbf{F}(M))$$

but this we have, since $M = \operatorname{Hom}_R(R, M)$, and $\mathbf{F}$ is a functor.

Exercise 5.1.45. (1) A priori, the map

$$\mathbf{F}_{R,M} : \mathrm{Hom}_R(R, M) \to \mathrm{Hom}_S(\mathbf{F}(R), \mathbf{F}(M))$$

is a map of *sets*. Check that it is, in fact, a map of R-modules (and therefore, so is t_M).

(2) Check that the resulting maps t_M form a natural transformation from $N \otimes_R -$ to $\mathbf{F}$

We note that so far, we did not use that $\mathbf{F}$ preserves colimits, the construction works in general. The assumption on $\mathbf{F}$ is employed via the following lemma:

Lemma 5.1.46. *Let $\mathbf{F}_1, \mathbf{F}_2 : R - \mathcal{M}od \to \mathcal{C}$ be two functors that commute with colimits, and let $t : \mathbf{F}_1 \to \mathbf{F}_2$ be a natural transformation. If $t_R : \mathbf{F}_1(R) \to \mathbf{F}_2(R)$ is an isomorphism or 0, then so is t.*

The proof will use presentations as above, which is now not a problem, because the map is already defined

Proof. Let M be an arbitrary module, and let $\oplus_\beta R \xrightarrow{f} \oplus_\alpha R \to M \to 0$ be a presentation. Since t is a natural transformation, we have a commutative diagram

$$\begin{tikzcd}
\mathbf{F}_1(\oplus_\beta R) \arrow[r, "\mathbf{F}_1(f)"] \arrow[d, "t_{\oplus_\beta R}"] & \mathbf{F}_1(\oplus_\alpha R) \arrow[r] \arrow[d, "t_{\oplus_\alpha R}"] & \mathbf{F}_1(M) \arrow[r] \arrow[d, "t_M"] & 0 \\
\mathbf{F}_2(\oplus_\beta R) \arrow[r, "\mathbf{F}_2(f)"] & \mathbf{F}_2(\oplus_\alpha R) \arrow[r] & \mathbf{F}_2(M) \arrow[r] & 0
\end{tikzcd} \tag{5.3}$$

Since both $\mathbf{F}_i$ commute with colimits, we obtain the following diagram, in which the rows are again exact (i.e., $\mathbf{F}_i(M)$ is the quotient by the image of $\mathbf{F}_i(f)$)

$$\begin{tikzcd}
\oplus_\beta \mathbf{F}_1(R) \arrow[r, "\mathbf{F}_1(f)"] \arrow[d, "\oplus_\beta t_R"] & \oplus_\alpha \mathbf{F}_1(R) \arrow[r] \arrow[d, "\oplus_\alpha t_R"] & \mathbf{F}_1(M) \arrow[r] \arrow[d, "t_M"] & 0 \\
\oplus_\beta \mathbf{F}_2(R) \arrow[r, "\mathbf{F}_2(f)"] & \oplus_\alpha \mathbf{F}_2(R) \arrow[r] & \mathbf{F}_2(M) \arrow[r] & 0
\end{tikzcd} \tag{5.4}$$

The fact that the vertical maps are as indicated follows from the exercise below. Given that, the two left vertical maps are isomorphisms or 0 if t_R has the same property, so the same holds for t_M. □

Exercise 5.1.47. Let $\mathbf{F}_1, \mathbf{F}_2 : \mathcal{C} \to \mathcal{D}$ be two functors, and let $t : \mathbf{F}_1 \to \mathbf{F}_2$ be a natural transformation. Let $\mathbf{I} : \mathcal{I} \to \mathcal{C}$ be a diagram in $\mathcal{C}$ with colimit $X = \varinjlim \mathbf{I}$, and assume that the natural map $\varinjlim(\mathbf{F}_1 \circ \mathbf{I}) \to \mathbf{F}_1(X)$ is an isomorphism (in particular, we assume the colimits exist). Show that t_X is the colimit map corresponding to the maps $\mathbf{F}_2(c_i) \circ t_{\mathbf{I}(i)}$, for i an object of $\mathcal{I}$, where $c_i : \mathbf{I}(i) \to X$ is the canonical map.

We summarise the proof of the theorem:

Proof of Theorem 5.1.43. Let $\mathbf{F} : R - \mathcal{M}od \to S - \mathcal{M}od$ be a functor, and let $N = \mathbf{F}(R)$, with the bimodule structure defined after the statement. The discussion prior to Lemma 5.1.46 constructs a natural map from $\mathbf{F}_N$ to $\mathbf{F}$, and the lemma itself shows that it is an isomorphism. The other direction is Prop. 5.1.40. □

Having settled the existence question, we may ask about uniqueness, i.e., can there be two bimodules N and N' for which $\mathbf{F}_N$ and $\mathbf{F}_{N'}$ are isomorphic? Of course, if $\mathbf{F}_N$ is isomorphic to some functor $\mathbf{F}$, then N is isomorphic to $N' = \mathbf{F}(R)$, and $\mathbf{F}$ is isomorphic to $\mathbf{F}_{N'}$. This suggests that we really need to consider the *categories* of functors and of modules:

Corollary 5.1.48. *Let R and S be rings. The association $N \mapsto \mathbf{F}_N = N \otimes_{R^\circ} -$ determines an equivalence of categories from the category $S \otimes R^\circ - \mathcal{M}od$ of $R - S$-bimodules to the category $\mathcal{R}ex(R, S)$ of colimit-preserving functors from $R - \mathcal{M}od$ to $S - \mathcal{M}od$.*

If T is a third ring, composition of functors corresponds to tensor products, i.e., $\mathbf{F}_M \circ \mathbf{F}_N$ is isomorphic to $\mathbf{F}_{M \otimes_S N}$, where N is an $R - S$-bimodule and M is an $S - T$-bimodule.

Proof. The inverse is given by evaluation on R, as in the proof of the Eilenberg–Watts Theorem. Then $\mathbf{F}_N(R) = N \otimes_R R = N$ (canonical isomorphism), and the isomorphism $cf \to \mathbf{F}_{\mathbf{F}(R)}$ was also given in the proof. The second part is an exercise. □

Corollary 5.1.49. *Let $f : R \to S$ be a map of rings, and let $\mathbf{F}_S$ be the corresponding base-change functor. Then* $\operatorname{End}(\mathbf{F}_S)$ *is isomorphic to the subring $Z_{S^\circ}(f(R))$ of S° consisting of elements commuting with the image $f(R)$ (this is the* centraliser *of $f(R)$ in S°). In particular, the endomorphism ring of the identity functor on $R - \mathcal{M}od$ is the centre $Z(R)$.*

$Z_{S^\circ}(f(R))$
centraliser

Proof. According to the Corollary above, $\operatorname{End}(\mathbf{F}_S)$ is isomorphic to the endomorphisms of S as an $R - S$-bimodule. We had already seen that its endomorphisms as an S module is S°, acting on the right. Since R acts on the right as well, the additional condition that the R-module is also preserved selects the elements of S° that commute with R. The last part follows by taking $S = R$ (and f the identity). □

End of lecture 18, May 25

5.2. **Morita equivalence.** We are now in position to answer a very natural question: to which extent does the category of modules $R - \mathcal{M}od$ determine the ring R? In other words, if two rings R and S have the same category of modules, what can be said of R and S? Of course, "the same" means equivalent in this case, so we arrive at the following definition:

Definition 5.2.1. Two rings R and S are *Morita equivalent* if their categories of modules are equivalent.

Morita equivalent

Let $\mathcal{C} = R - \mathcal{M}od$ and $\mathcal{D} = S - \mathcal{M}od$. If $\mathbf{F} : \mathcal{C} \to \mathcal{D}$ is an equivalence, it definitely preserves colimits, and so has the form $\mathbf{F}_P$ for some $S \otimes R^\circ$-module N. Since $\mathbf{F}$ is an equivalence, there is also a functor $\mathbf{G} : \mathcal{D} \to \mathcal{C}$, and suitable compositions between their compositions and the identity. Translating to modules, and using Corollary 5.1.48, we obtain

Corollary 5.2.2. *$R - \mathcal{M}od$ is equivalent to $S - \mathcal{M}od$ if and only if there are an $S \otimes R^\circ$-module P, an $R \otimes S^\circ$-module Q, and isomorphisms $f : Q \otimes_S P \xrightarrow{\sim} R$ of $R \otimes R^\circ$-modules, and $g : P \otimes_R Q \xrightarrow{\sim} S$ of $S \otimes S^\circ$-modules*

The typical example is as follows:

Example 5.2.3. Let R be a ring, $P \neq 0$ a free module of finite rank over R°, and $S = \mathrm{End}_R(P)$. Then R and S are Morita equivalent. Indeed, let $Q = \mathrm{Hom}_R(P, R)$. The action of S on P determines an action of S° on Q, and R acts by acting on the image from the left. We have the evaluation map $f : Q \otimes_S P \to R$, and the isomorphism $g : P \otimes_R Q \to S$ from Prop. 3.2.1. Although f is surjective, at the moment it is not obvious that this is an isomorphism. □

To show that the map f above is an isomorphism, we note an additional property: Assume we are given modules P, Q and maps $f : Q \otimes_S P \to R$ and $g : P \otimes_R Q \to S$ as in Corollary 5.2.2, except that f and g are not required to be isomorphisms. There are, then, two maps $P \otimes_R Q \otimes_S P \to P$ (of $S \otimes R^\circ$-modules), namely $g \otimes 1$ and $1 \otimes f$. Likewise, there are two maps $Q \otimes_S P \otimes_R Q \to Q$, given by $f \otimes 1$ and $1 \otimes g$.

Exercise 5.2.4. Show that in the situation of Example 5.2.3, the two maps $g \otimes 1$ and $1 \otimes f$ are equal, and also the two maps $1 \otimes g$ and $f \otimes 1$ are equal.

This suggests the following definition:

Morita context

Definition 5.2.5. Let R and S be two rings. A *Morita context* for R, S consists of an $S \otimes R^\circ$-module P, an $R \otimes S^\circ$-module Q, and bi-module maps $f : Q \otimes_S P \to R$ and $g : P \otimes_R Q \to S$, such that $g \otimes 1 = 1 \otimes f$ and $1 \otimes g = f \otimes 1$.

Exercise 5.2.6. Let (P, Q, f, g) be data as in a Morita context, not necessarily satisfying the conditions. Let $M = \{\left(\begin{smallmatrix} r & q \\ p & s \end{smallmatrix}\right) \mid r \in R, p \in P, q \in Q, s \in S\}$. We may define an operation on M via matrix multiplication, where the "products" are given by $qp = f(q \otimes p)$ and $pq = g(p \otimes q)$ (and by the various module structures). Show that this operation is associative precisely if these data form a Morita context (and thus, in this case M is a ring).

To complete our typical example, we show:

Proposition 5.2.7. *Let (P, Q, f, g) be a Morita context. If f is surjective, then it is injective. Hence, if both f and g are surjective, then the Morita context determines a Morita equivalence between R and S.*

Proof. Since f is surjective, there is $x \in Q \otimes_S P$ with $f(x) = 1$. Assume that $f(y) = 0$, and consider $x \otimes y \in Q \otimes P \otimes Q \otimes P$. We have

$$y = 1 \otimes y = f(x) \otimes y = (f \otimes 1 \otimes 1)(x \otimes y) = (1 \otimes g \otimes 1)(x \otimes y) = \\ (1 \otimes 1 \otimes f)(x \otimes y) = x \otimes f(y) = x \otimes 0 = 0$$

□

This completes the example, but raises the question: when does an equivalence of categories determine a Morita context? We may formulate the question on the level of abstract categories, as follows.

Assume that $\mathbf{F} : \mathcal{C} \to \mathcal{D}$ is an equivalence of categories. This is witnessed by the existence of a functor $\mathbf{G} : \mathcal{D} \to \mathcal{C}$ and two isomorphisms, $\alpha : \mathbf{G} \circ \mathbf{F} \to \mathbf{1}_{\mathcal{C}}$ and $\beta : \mathbf{F} \circ \mathbf{G} \to \mathbf{1}_{\mathcal{D}}$. If γ is an automorphism of $\mathbf{1}_{\mathcal{C}}$, we may replace α by $\alpha' = \gamma \circ \alpha$, and obtain new such data. Conversely, every other such α' is obtained via some $\gamma \in \mathrm{Aut}(\mathbf{1}_{\mathcal{C}})$, namely, $\gamma = \alpha' \circ \alpha^{-1}$. We may remove the arbitrariness in the choice of α as follows: For any object X of $\mathcal{C}$, $\alpha_X : \mathbf{G} \circ \mathbf{F}(X) \to X$ is a morphism in $\mathcal{C}$, so we may apply $\mathbf{F}$ to obtain an isomorphism $\mathbf{F}(\alpha_X) : \mathbf{F}(\mathbf{G}(\mathbf{F}(X))) \to \mathbf{F}(X)$. On the other hand, $\beta_{\mathbf{F}(X)}$ is another such isomorphism, and so we may ask if they are equal. It is easy to see, at least, that the collection $\mathbf{F}(\alpha_X)$ forms a natural isomorphism from $\mathbf{F} \circ \mathbf{F} \circ \mathbf{F}$ to $\mathbf{F}$, which we denote by $\mathbf{F}(\alpha)$. Likewise, we have a natural isomorphism $\mathbf{G}(\beta) : \mathbf{G} \circ \mathbf{F} \circ \mathbf{G} \to \mathbf{G}$, and we may ask if it coincides with $\alpha_{\mathbf{G}}$.

Proposition 5.2.8. *Assume that $\mathbf{F} : \mathcal{C} \to \mathcal{D}$ is an equivalence of categories, with quasi-inverse $\mathbf{G} : \mathcal{D} \to \mathcal{C}$.*

(1) *For every isomorphism $\beta : \mathbf{F} \circ \mathbf{G} \to \mathbf{1}_{\mathcal{D}}$ there is a unique isomorphism $\alpha : \mathbf{G} \circ \mathbf{F} \to \mathbf{1}_{\mathcal{C}}$ such that $\mathbf{F}(\alpha) = \beta_{\mathbf{F}}$.*

(2) *If α and β are as above, then $\mathbf{G}(\beta) = \alpha_{\mathbf{G}}$*

Proof. (1) By assumption, there is some isomorphism $\alpha' : \mathbf{G} \circ \mathbf{F} \to \mathbf{1}_{\mathcal{C}}$. For every object X of $\mathcal{C}$, $\beta_{\mathbf{F}(X)} \circ \mathbf{F}({\alpha'_X}^{-1})$ is an automorphism of $\mathbf{F}(X)$, so is equal to $\mathbf{F}(\gamma_X)$ for a unique automorphism γ_X of X (since $\mathbf{F}$ is bijective on morphisms). The uniqueness easily implies that this collection forms a natural automorphism γ of the identity. Let $\alpha = \gamma \circ \alpha'$. Then we have

$$\mathbf{F}(\alpha) = \mathbf{F}(\gamma) \circ \mathbf{F}(\alpha') = \beta_{\mathbf{F}} \circ \mathbf{F}({\alpha'}^{-1}) \circ \mathbf{F}(\alpha') = \beta_{\mathbf{F}}$$

as required.

(2) Since $\mathbf{F}$ is injective on morphisms, it suffices to show that $\mathbf{F}(\mathbf{G}(\beta)) = \mathbf{F}(\alpha_{\mathbf{G}})$. By assumption, $\mathbf{F}(\alpha_{\mathbf{G}}) = \beta_{\mathbf{F} \circ \mathbf{G}}$, so it remains to show that $\mathbf{F}(\mathbf{G}(\beta)) = \beta_{\mathbf{F} \circ \mathbf{G}}$, which follows immediately from the more general formula $\mathbf{F}(\mathbf{G}(t)) = \beta_Y^{-1} \circ t \circ \beta_X$ for any morphism $t : X \to Y$ in $\mathcal{D}$. □

Exercise 5.2.9. Deduce that any equivalence between categories of modules determines a Morita context.

Our goal now is to describe the information required for an equivalence in a minimal way. The first observation is that all the data is determined by one module:

Proposition 5.2.10. *Assume that (P, Q, f, g) is a Morita context that determines an equivalence $R - \mathcal{M}od \to S - \mathcal{M}od$.*

(1) For any S-module M, there is an R-module isomorphism $\alpha_M : Q \otimes_S M \xrightarrow{\sim} \operatorname{Hom}_S(P, M)$, natural[5] *in M.*

(2) There are R-module isomorphisms $g^\vee : Q \xrightarrow{\sim} \operatorname{Hom}_S(P, S)$ and $h : R \xrightarrow{\sim} \operatorname{End}_S(P)$, such that the diagrams

$$\begin{tikzcd}
Q \otimes_S P \arrow[r, "f"] \arrow[d, "g^\vee \otimes 1"] & R \arrow[d, "h"] \\
\operatorname{Hom}_S(P, S) \otimes_S P \arrow[r, "t"] & \operatorname{End}_S(P)
\end{tikzcd} \tag{5.5}$$

and

$$\begin{tikzcd}
P \otimes_R Q \arrow[d, "1 \otimes g^\vee"] \arrow[dr, "g"] & \\
P \otimes_R \operatorname{Hom}_S(P, S) \arrow[r, "\mathrm{ev}"] & S
\end{tikzcd} \tag{5.6}$$

commute. Here, ev *is the evaluation, and $t(\varphi \otimes p)(p') = \varphi(p')p$*

Proof. (1) To define such a map is equivalent to defining an S-module map $P \otimes_R Q \otimes_S M \to M$, which we take to be $g \otimes 1$. To define a map in the other direction, we consider the S-module map

$$P \otimes_R \operatorname{Hom}_S(P, M) \xrightarrow{\mathrm{ev}} M \xrightarrow{g^{-1} \otimes 1} P \otimes_R Q \otimes_S M$$

Applying $Q \otimes_S -$ and composing with f, we obtain a map $\operatorname{Hom}_S(P, M) \to Q \otimes_S M$. The fact that the two maps are inverse to each other follows from the context condition.

(2) The isomorphisms are obtained from the previous claim by taking $M = S$ and $M = P$, respectively, and using f to identify $Q \otimes_S P$ with R. We may then enhance the first diagram to

$$\begin{tikzcd}
Q \otimes_S P \arrow[r, "f"] \arrow[d, "g^\vee \otimes 1"] \arrow[dr, "\alpha_P"] & R \arrow[d, "h"] \\
\operatorname{Hom}_S(P, S) \otimes_S P \arrow[r, "t"] & \operatorname{End}_S(P)
\end{tikzcd}$$

with the top triangle commuting by definition, so we have to show that $\alpha_S(q)(p')p = \alpha_P(q \otimes p)(p')$ and that $\alpha_S(q)(p) = g(p \otimes q)$, both of which follow directly from the construction. □

[5] This means that the collection α forms a natural isomorphism

Having seen that the full data of the Morita equivalence is determined by P as an S-module, we would like to understand what are the properties of such a P. By the first part of the proposition, the functor $M \mapsto \mathrm{Hom}_S(P, M)$ determines an equivalence of categories, so in particular, it preserves all possible limits. It is easy to see that such a functor always commutes with all inverse limits (for any P), but preservation of colimits is a non-trivial condition. We had seen in Remark 5.1.39 that this condition is equivalent to two special cases: preservation of cokernels (quotients), and of coproducts. The first condition is interesting on its own:

Definition 5.2.11. An S-module P is a *projective module* if the functor $\mathrm{Hom}_S(P, -)$ is *exact*, i.e., for all surjective S-module maps $A \to B$, the induced map $\mathrm{Hom}_S(P, A) \to \mathrm{Hom}_S(P, B)$ is surjective as well.

projective module

We note that since in the category of modules, finite products coincide with finite coproducts, the above condition is equivalent to preservation of all finite colimits. In more concrete terms, the condition means that given any surjective map $p : A \to B$, and any $t : P \to B$, there is a map $\tilde{t} : P \to A$ such that $p \circ \tilde{t} = t$.

Exercise 5.2.12.
1. Show that any free module is projective
2. Show that a module P is projective if and only if there is a free module F and an isomorphism $F \xrightarrow{\sim} P \oplus Q$ for some module Q
3. Show that if V is a vector-space over a field k of finite dimension greater than 1, then V is projective but not free over $\mathrm{End}_k(V)$

End of lecture 19, Jun 8

The other part of the condition on the preservation of limits is about preservation of (infinite) direct sums. In concrete terms, assume that A_α is a family of modules. A map $P \to A_i$ can be composed with the inclusion to induce a map into $\oplus_\alpha A_\alpha$. Taking maps, this describes the map $\oplus_\alpha \mathrm{Hom}(P, A_\alpha) \to \mathrm{Hom}(P, \oplus_\alpha A_\alpha)$. For this to be an isomorphism means that every map from P to the direct sum is a finite sum of maps to the components. This need not be the case for a general P: for example, the identity map from $\oplus_\alpha A_\alpha$ to itself is not of this form if the sum is infinite. However, if P is finitely generated, then the map is determined by its image on the generators, which lies in a finite sum. This proves the first part of:

Proposition 5.2.13. *If P is a finitely generated R-module, then for every family A_α of R-modules, the map $\oplus_\alpha \mathrm{Hom}_R(P, A_\alpha) \to \mathrm{Hom}_R(P, \oplus_\alpha A_\alpha)$ is an isomorphism. If P is projective, the converse is also true.*

Proof. The first part was explained above. Assume that P is projective, and write $P \oplus Q = F = R^{\oplus \alpha}$, as in Exercise 5.2.12. By assumption, the map from P to F factors through a finite sum F_0. The projection from F_0 to P gives a finite set of generators. □

We need one last condition: We have been repeatedly using the fact that every object has a surjective map from some direct sum of copies of the ring (for example, in the last proof). In categorical terms, we say that R generates its category of module, in the following sense.

set of generators

Definition 5.2.14. A *set of generators* for a category $\mathcal{C}$ is a set S of objects of $\mathcal{C}$ such that whenever $f, g : X \to Y$ are two distinct maps, there is a map $p : P \to X$, with $P \in S$, such that $f \circ p \neq g \circ p$.

Intuitively, maps $p : P \to X$ can be thought of as "generalised points of X of type P". From this point of view, the definition says that every object has "sufficiently many" P-points ($P \in S$) to distinguish between morphisms.

Example 5.2.15. Any non-empty set is a generator for the category of sets. □

Example 5.2.16. Any non-zero free module is a generator for the category of R-modules □

epimorphism

Exercise 5.2.17. A morphism $p : P \to X$ is an *epimorphism* if for any $f, g : X \to Y$, if $f \circ p = g \circ p$ then $f = g$.

(1) Assume that $\mathcal{C}$ has coproducts. Show that S is a set of generators if and only if for any object X there is a coproduct P of objects of S, and an epimorphism $p : P \to X$
(2) Show that in the category of R-modules, a map is an epimorphism if and only if it is surjective.

Example 5.2.18. If S is a set of simple R-modules that has a representative from each isomorphism class, then S generates the opposite of the category of JSS R-modules (i.e., it cogenerates the category of JSS R-modules) □

Proposition 5.2.19. *Assume P is a projective R-module. Then it generates $R - \mathcal{M}od$ if and only if for any non-zero R-module M there is a non-zero map $P \to M$.*

Proof. If P generates $R - \mathcal{M}od$, there will be a surjective map from a direct sum of copies of P onto M. One of the components must be non-zero (this direction does not use projectivity).

In the other direction, to show that P generates it suffices to show that if $f : X \to Y$ is non-zero, then $f \circ p \neq 0$ for some $p : P \to X$ (take the difference $f - g$ in the definition). By assumption, there is a non-zero map $p_0 : P \to M = f(X)$ to the image of f. Since P is projective, it can be lifted to a (non-zero) map $p : P \to X$ □

We may summarise what we learned so far as follows:

Corollary 5.2.20. *Assume that $\mathbf{F} : R - \mathcal{M}od \to S - \mathcal{M}od$ is an equivalence of categories. Then $P = \mathbf{F}(R)$ is a finitely generated projective generator, R is isomorphic to $\mathrm{End}_S(P)^\circ$, and an inverse to $\mathbf{F}$ is given by $M \mapsto \mathrm{Hom}_S(P, M)$.*

We now claim the converse:

Theorem 5.2.21. *Assume that P is a finitely generated projective generator of the category $S-\mathcal{M}od$. Then the functor $M \mapsto \mathrm{Hom}_S(P,M)$ determines an equivalence of categories $S-\mathcal{M}od \to R-\mathcal{M}od$, where $R = \mathrm{End}_S(P)^\circ$. An inverse is given by $N \mapsto P \otimes_R N$.*

Proof. Let $Q = \mathrm{Hom}_S(P,S)$, let $g : P \otimes_R Q \to S$ be the evalution, and let $f : Q \otimes_S P \to R$ be given by $f(\varphi \otimes p)(p') = \varphi(p')p$ (compare Prop. 5.2.10). We have seen that (P,Q,f,g) is a Morita context, so by Prop. 5.2.7, to prove the claim it suffices to prove that f and g are surjective, i.e., that 1 is in the image.

Since P is a generator, there is a surjective map from a direct sum of copies of P to S. Since $1 \in S$ is in the image of a finite sub-sum, we may assume that the sum itself is finite. Hence, we have $q_1,\dots,q_n \in Q$, and $p_1,\dots,p_n \in P$, such that $g(\sum p_i \otimes q_i) = \sum q_i(p_i) = 1$, as required.

For f, write $P \oplus P' = F$, where F is a finitely-generated free S-module. Then $G = \mathrm{Hom}_S(F,S) = \mathrm{Hom}_S(P,S) \oplus \mathrm{Hom}_S(P',S) = Q \oplus Q'$. We have the analogous map

$$\tilde{f} : G \otimes F \to \mathrm{End}_S(F) = \mathrm{End}(P) \oplus \mathrm{End}(P') \oplus \mathrm{Hom}(P,P') \oplus \mathrm{Hom}(P',P)$$

and $\tilde{f}(G \otimes F) \cap \mathrm{End}(P)$ is the image of f. But we had seen in Example 5.2.3 that $\tilde{f}$ is surjective, hence so is f. □

Example 5.2.22. Assume A is a finite-dimensional algebra over a field. If e is a non-trivial idempotent (i.e., $e^2 = e$, $e \neq 0,1$), then we may write A as a direct sum $A = Ae \oplus A(1-e)$ of left A-modules. By dimension, there is a finite maximal set of idempotents e_i with $e_ie_j = 0$ for $i \neq j$, and a corresponding decomposition $A = \oplus_i Ae_i$. A is called *basic* if all these summands are non-isomorphic.

Let $e_{i_1},\dots,e_{i_k}$ be a subset corresponding to representative of isomorphism classes, and let $e = e_{i_1} + \dots + e_{i_k}$. Let $A' = eAe$. Then A' is a basic algebra. We claim that A' is Morita equivalent to A.

Let $P = Ae_{i_1} \oplus \dots \oplus Ae_{i_k}$. We claim that P is a finitely generated projective generator of $A-\mathcal{M}od$. The fact that it is finitely generated is obvious, it is projective since it is a direct summand of the free module A, and it is a generator since all isomorphism classes of the summands of A are represented. Finally, a direct computation shows that $\mathrm{End}_A(P) = A'$. □

Exercise 5.2.23. Compute the above example for the case $A = M_n(k)$, k a field.

We will see more examples in §6.

5.3. **Morita invariant properties.**

We now go back to our original question: What properties of a ring are preserved under Morita equivalence? In other words, what properties of a ring are actually properties of its category of modules? We essentially saw the following:

Corollary 5.3.1. *If R and S are Morita equivalent, then $Z(R) \xrightarrow{\sim} Z(S)$. In particular, Morita equivalent commutative rings are isomorphic.*

This corollary explains why the notion of Morita equivalence is not important in commutative algebra. On the other hand, a commutative ring can be Morita equivalent to a non-commutative one (for example, k is equivalent to $\mathrm{End}_k(V)$), so the property "being commutative" is not Morita invariant.

Proof. We saw in Corollary 5.1.49 that the centre is the ring of endomorphisms of the identity functor on the category of modules (which is obviously preserved under equivalence) □

An easy corollary of the description in terms of contexts is:

Corollary 5.3.2. *If R and S are Morita equivalent, then so are R° and S°.*

Exercise 5.3.3. Prove it

To go further, it is useful first to note which subcategories of the category of modules are preserved under equivalence.

Proposition 5.3.4. *Let $\mathbf{F} : R - \mathcal{M}od \to S - \mathcal{M}od$ be an equivalence of categories. Then $\mathbf{F}$ restricts to an equivalence of the following full subcategories of R-modules with their S-module counterparts*

1. *$R - \mathcal{P}roj$, the projective R-modules*
2. *$R - \mathcal{V}ec$, the finitely generated projective R-modules*
3. *$R - \mathcal{C}oh$, finitely generated modules R-modules*
4. *The simple, semisimple and JSS R-modules*
5. *Artinian module, Noetherian modules*

Proof.
1. The definition of a projective module is categorical
2. These have a categorical description by Prop. 5.2.13.
3. We know that the equivalence is given by tensoring with a finitely generated module, which preserves finite generation
4. A module M is simple if and only if every morphism from another module N is either 0 or an epimorphism, so this is a categorical condition (a morphism is 0 if and only if it factors through the 0 module, which is the initial and final object). Semisimplicity is defined categorically, and M is JSS if and only if it has an injective map into a direct sum of simple modules.
5. Directly from the definition □

We immediately obtain:

Corollary 5.3.5. *The following properties of a ring are preserved under Morita equivalence*

1. *Semisimple*
2. *Artinian, Noetherian, finite length*
3. *JSS*

Proof.
1. This is equivalent to all modules being semisimple
2. This is a property shared by all finitely generated modules (Prop. 4.3.4), so is preserved by the last proposition.

(3) This is shared by the projective finitely generated modules: Each such module is a direct summand of a free module, so is JSS by Prop. 4.1.7 □

End of lecture 20, Jun 14

6. Division algebras

6.1. **Central simple algebras.** Our initial goal is to study the class of finite dimensional division algebras over a fixed field k. We had seen these algebra in the context of the group algebra (Prop. 2.5.1), but they can also be viewed as generalizations of finite field extensions, and are important in the study of the base field k. Taking this point of view, we note that the centre $Z(D)$ of a division k-algebra is a finite dimensional subfield of D containing k, i.e., a finite field extension of k. Since such extensions are studied in Galois theory, it makes sense to consider separately the case where the extension is trivial.

Definition 6.1.1. A *central division algebra* over a field k is a division algebra of finite dimension over k, whose centre is k.

central division algebra

One case where this studying such extensions is easy is when k is algebraically closed: in this case, k is the only such algebra. So one way to try to understand division algebras over a general field is to base-change to its algebraic closure, and descend from there. Of course, this base-change will no longer be a division algebra, so we would like to slightly extend our class of algebras, so that it is closed under tensor products.

We fix the field k. All our rings in this section will be k-algebras, and tensor products will be over k, even if this is suppressed from notation. Since we wish to keep track of the centre, we first compute:

Proposition 6.1.2. *For any k-algebras A and B, $Z(A \otimes B) = Z(A) \otimes Z(B)$.*

Proof. We first note that for any algebra C, the centre is the kernel of the linear map $c_C : C \to \mathrm{End}(C)$ sending c to $a \mapsto ca - ac$. In our case, since $Z(B)$ is the kernel of c_B and k is a field, we have that $A \otimes Z(B)$ is the kernel of $1 \otimes c_B : A \otimes B \to A \otimes \mathrm{End}(B) = \mathrm{Hom}(B, A \otimes B)$. On the other, hand, we have a restriction map $r : \mathrm{End}(A \otimes B) \to \mathrm{Hom}(B, A \otimes B)$, and a direct computation shows that $1 \otimes c_B = r \circ c_{A\otimes B}$. Hence, $Z(A \otimes B) = \ker(c_{A\otimes B}) \subseteq A \otimes Z(B)$. Applying the same result again (with A playing the role of B), we get $Z(A \otimes B) = Z(Z(A \otimes B)) \subseteq Z(A \otimes Z(B)) \subseteq Z(A) \otimes Z(B)$. The other inclusion is obvious. □

Remark 6.1.3. In the proof we used the following: If $U \subseteq V$ are vector-spaces over k, and W is an additional vector-space, then the natural map $W \otimes_k U \to W \otimes V$ is injective. This follows from Exercise 3.1.5. The proof, more generally, work when k is a ring, provided at least one of the algebras is flat over k. □

The proposition implies, in particular, that the property of being central is preserved when extending the field (i.e., making a base-change by a field extension). The property of being a division ring is not preserved, but we will see that the somewhat weaker property of being *simple* (i.e., having no two-sided ideals) is preserved. Thus it makes sense to consider:

central simple algebra
CSA

Definition 6.1.4. A *central simple algebra* (*CSA*) over the field k is a finite dimensional simple algebra over k whose centre is k.

Example 6.1.5. Over an algebraically closed field k, every CSA is isomorphic to $\mathrm{End}_k(V)$ for some finite dimensional vector space V. This follows from Prop. 2.5.1, along with the fact that a simple ring cannot be a non-trivial Cartesian product □

The argument in the last example shows, more generally, that each CSA is isomorphic to $M_n(D)$ for a unique central division algebra D and a unique n (by dimension).

We had seen that the base change of a central division algebra to a field extension is a central over that field. We promised that it will be simple, and one could ask more generally, what happens to a CSA after base change. The answer, as well as additional properties, follows from the following statement:

Proposition 6.1.6. *Let A be a CSA over k. Then for every k-algebra B, the functor $M \mapsto A \otimes_k M$ is an equivalence of categories from $B \otimes B^\circ - \mathcal{M}od$ to $A \otimes B \otimes (A \otimes B)^\circ - \mathcal{M}od$.*

Here, the bi-module structure on $A \otimes M$ is the obvious one, with $B \otimes B^\circ$ acting on M and $A \otimes A^\circ$ acting on A.

Corollary 6.1.7. *Let A be a CSA over k*

1. *(1) For any algebra B over k, the map $I \mapsto I \cap A$ determines a bijection between two sided ideals in $A \otimes B$ and two sided ideals in B, with inverse $J \mapsto A \otimes J$.*
2. *(2) If l is a field extension of k then $A \otimes l$ is a CSA over l*
3. *(3) If B is another CSA over k, then so is $A \otimes B$*

Proof. For the last two parts, Proposition 6.1.2 shows that the centre is correct, and Exercise 3.1.4 shows that the dimension is finite. Simplicity follows, in each case, from the first part. So it remains to prove the first part.

We note that a two-sided ideal in a ring C is the same as a left submodule of C, viewed as a module over $C \otimes C^\circ$. By Proposition 6.1.6, such submodules for B are in bijective correspondence with those over $A \otimes B \otimes A \otimes B^\circ$. □

The main step in the proof of Proposition 6.1.6 is the following fact, which is sometimes taken as the definition of CSAs:

Lemma 6.1.8. *Let A be a finite dimensional simple k-algebra. Then A is a CSA over k if and only if the map $t : A \otimes A^\circ \to \mathrm{End}_k(A)$ given by $a \otimes b \mapsto (c \mapsto acb)$ is an isomorphism.*

Proof. If $a \in Z(A)$ is not in k, then $a \otimes 1 - 1 \otimes a$ is a non-zero element in the kernel of t. For the converse, by the basic Structure Theorem, A is isomorphic to $M_n(D)$ for some division ring D. Since $M_n(D)$ is isomorphic to $M_n(k) \otimes D$, and $\mathrm{End}_k(M_n(D))$ is isomorphic to $\mathrm{End}_k(M_n(k)) \otimes \mathrm{End}_k(D)$, we are reduced to the case $A = D$ is a division algebra (the case $A = M_n(k)$ is easy).

Assume that $x = \sum_{i=1}^{n} d_i \otimes e_i$ is in the kernel of t, with n minimal. Multiplying by inverses, we may assume that $d_1 = e_1 = 1$. Since the kernel of t is a two-sided ideal, each commutator $[x, d \otimes e]$ is also in the kernel. But this commutator is the sum of less tensors, so is 0 by minimality. Since the tensors span $D \otimes D^\circ$, this shows that x is in the centre. Prop. 6.1.2 shows that all d_i and e_i belong to the centre of D. By assumption, $x \in k = k \otimes k$, and thus $x = 0$. This shows that t is injective, and by dimension, bijective. □

Proof of 6.1.6. Set $S = B \otimes B^\circ$, and let $Q = A \otimes S$. Then for all S-modules M, $A \otimes M = Q \otimes_S M$. Setting $P = \mathrm{Hom}_S(Q, S) = \mathrm{Hom}_k(A, k) \otimes S$, we see that P is free of finite rank (equal to the dimension of A over k), and is thus a projective generator. Theorem 5.2.21 thus applies, so that $M \otimes_k - = Q \otimes_S - = \mathrm{Hom}_S(P, -)$ is an equivalence from $S - \mathcal{M}od$ to the category of R-modules, where $R = \mathrm{End}_S(P)$.

Now,

$$R = \mathrm{End}_S(P) = \mathrm{End}_S(A^\vee \otimes S) = \mathrm{End}_k(A) \otimes S = \mathrm{End}_k(A) \otimes B \otimes B^\circ$$

By Lemma 6.1.8, we have $\mathrm{End}_k(A) = A \otimes A^\circ$, so we are done. □

We now draw some conclusions:

Corollary 6.1.9. *The dimension of every CSA as a vector space over k is a square of an integer*

Proof. The dimension is preserved under base-change, so we may assume the base is algebraically closed. There, every CSA is a matrix algebra. □

Before proceeding, we note that Morita equivalence behaves well with respect to tensor product:

Proposition 6.1.10. *If A is Morita equivalent to A' and B is Morita equivalent to B', then $A \otimes B$ is Morita equivalent to $A' \otimes B'$*

Proof. It suffices to check for A and B separately, so we may assume $B = B'$. Assume that $\mathbf{F} : A - \mathcal{M}od \to A' - \mathcal{M}od$ is an equivalence. Any $A \otimes B$-module is, in particular, an A module, so $\mathbf{F}(A)$ is an A'-module. The action of B on M is given by maps in $A - \mathcal{M}od$, so induces a commuting B-module structure on $\mathbf{F}(A)$. It is easy to check that this makes $\mathbf{F}$ a functor from $A \otimes B - \mathcal{M}od$ to $A' \otimes B - \mathcal{M}od$. An inverse of $\mathbf{F}$ on $A - \mathcal{M}od$ will also be an inverse on $A' \otimes B - \mathcal{M}od$. □

We now strengthen the tie between CSAs and division algebras.

Proposition 6.1.11. *Let A and B be CSAs over k. The following are equivalent:*

(1) A and B are Morita equivalent
(2) There is a division algebra D such that $A \xrightarrow{\sim} M_n(D)$ and $B \xrightarrow{\sim} M_{m\times m}(D)$ for some n, m
(3) $M_{m\times m}(A) \xrightarrow{\sim} M_n(B)$ For some n, m

In particular, any Morita equivalent division algebras over k are isomorphic.

Proof.

(1) $\implies$ (2)**:** We already mentioned that the statement is true for different division algebras, so we need to show that Morita equivalent division algebras D and D' are isomorphic. Since the category of modules over a division ring has a unique simple object, an equivalence must take D to (an object isomorphic to) D', hence $D^\circ = \mathrm{End}_D(D) = \mathrm{End}_{D'}(D') = D'^\circ$.

(2) $\implies$ (3)**:** is clear since $M_n(M_{m\times m}(D)) = M_{nm\times nm}(D)$

(3) $\implies$ (1)**:** We saw that $M_{m\times m}(A)$ is Morita equivalent to A □

The proposition implies that studying isomorphism classes of central division algebras over k amounts to studying Morita equivalence classes of CSA over k. If A and B are such CSAs, Cor. 6.1.7 shows that $A\otimes B$ is also a CSA, and Prop. 6.1.10 shows that the Morita class $[A\otimes B]$ of $A\otimes B$ depends only on the Morita classes of A and B. Thus we have a well defined operation $[A]\otimes[B] := [A\otimes B]$. This operation is clearly associative and commutative, and since $A\otimes k \xrightarrow{\sim} A$, $[k]$ serves as a unit. Finally, Lemma 6.1.8 shows that $[A^\circ]$ is the inverse of $[A]$. Thus, the collection of Morita classes of CSAs over k forms a group.

$[A \otimes B]$

Brauer group
$\mathrm{Br}(k)$

Definition 6.1.12. The *Brauer group* $\mathrm{Br}(k)$ of a field k is the group of Morita equivalence classes of central simple algebras over k, with operation induced by tensor product.

We note again that, as a set, $\mathrm{Br}(k)$ is canonically isomorphic to the set of isomorphism classes of central division algebras over k, but the description of the operation is more complicated in this case.

End of lecture 21, Jun 15

6.2. **Some Brauer groups.** In this subsection, we prove a few important results, which will enable us (among other things) to classify finite division algebras. But first we consider the case $k = \mathbb{R}$.

Quaternions algebra

The *Quaternions algebra* is the algebra $\mathbb{H}$ over $\mathbb{R}$ spanned by $1, i, j, k$ as a vector space, with product determined by the relations $i^2 = j^2 = k^2 = -1$, $ij = k = -ji$.

Theorem 6.2.1 (Frobenius's Theorem)**.** *The only central division algebras over $\mathbb{R}$ are $\mathbb{R}$ and $\mathbb{H}$.*

Proof. Let D be a non-commutative central division algebra over $\mathbb{R}$. If $x \in D \setminus \mathbb{R}$, then $\mathbb{R}(x) \subseteq D$ is a commutative field, hence is isomorphic to $\mathbb{C}$. We may therefore assume that $\mathbb{C} = \mathbb{R}(i)$, with $i^2 = -1$.

Let $\varphi : D \to D$ be given by $\varphi(a) = i^{-1}ai$. It is $\mathbb{C}$-linear, with φ^2 the identity. Therefore, it has two eigenspaces, D_+ and D_-, corresponding to the eigenvalues 1 and -1, respectively. Since i is not in the centre, D_- is non-zero, and if $y \in D_-$ is non-zero, then $z \mapsto y^{-1}z$ is an $\mathbb{R}$-linear isomorphism from D_- to D_+.

We note that D_+ contains $\mathbb{C}$. If $z \in D_+ \setminus \mathbb{C}$, then $\mathbb{C}(z) \subseteq D$ is a commutative field extension of $\mathbb{C}$, a contradiction. Hence, D_+ and D_- are 1-dimensional over $\mathbb{C}$. If $j \in D_-$ is non-zero, then $\mathbb{R}(j)$ is a proper field extension of $\mathbb{R}$, hence is isomorphic to $\mathbb{C}$, but distinct from the copy of $\mathbb{C}$ we already have. We may assume that $j^2 = -1$. We define $k = ij$, and deduce the identities of $\mathbb{H}$ from the fact that j is an eigenvector. □

Exercise 6.2.2. Find an explicit isomorphism $\mathbb{C} \otimes_{\mathbb{R}} \mathbb{H} \to M_2(\mathbb{C})$

6.2.3. *Automorphisms of CSAs.* If $u \in R$ is a unit of an algebra R, then $r \mapsto uru^{-1}$ is an automorphism. An *inner automorphism* is an automorphism of this form.

inner automorphism

Theorem 6.2.4 (Skolem–Noether). *Let $R \subset S$ be a simple subalgebra of a CSA S, and let $f : R \to S$ be a homomorphism. Then there is a unit $u \in S$ such that $f(r) = uru^{-1}$ for all $r \in R$. In particular, every automorphism of S is inner.*

Proof. Let S' be S, viewed as an R-module via f. We view both S and S' as modules over $R \otimes S^\circ$. Clearly, S and S' are simple, and since $R \otimes S^\circ$ is simple (hence primitive), S and S' are isomorphic. Let $t : S \to S'$ be an isomorphism. Then for all $r \in R$ and $x, s \in S$, we have $t(rxs) = f(r)t(x)s$. For $s = 1$ we get $t(rx) = f(r)t(x)$ and for $s = r, r = 1$ we get $t(xr) = t(x)r$, so setting $u = t(1)$, we get $f(r)u = ur$. Setting $v = t^{-1}(1)$, we find $1 = t(v) = t(1v) = uv$ and similarly in the other direction, so $v = u^{-1}$. Hence $f(r) = uru^{-1}$. □

Exercise 6.2.5. Show that if $a, b \in R$ (a CSA over k) have the same minimal polynomial over k, then $a = ubu^{-1}$ for some unit

6.2.6. *The centraliser theorem.*

Proposition 6.2.7. *Let S be a CSA over k, and let $R \subseteq S$ be a simple subalgebra. Let $T = Z_S(R)$ be the centraliser of R in S.*

(1) T° is Morita equivalent to $R \otimes S^\circ$
(2) T is simple
(3) $[S : k] = [R : k][T : k]$

Proof. (1) Since S° is a CSA and R is simple, $R \otimes S^\circ$ is simple. Hence, every non-zero module over it is a projective generator. Consider the $R \otimes S^\circ$-module S. It is clearly finitely generated, so determines a Morita equivalence with $\mathrm{End}_{R\otimes S^\circ}(S)^\circ$. We had already computed (in Cor. 5.1.49) that this endomorphism ring is T.

(2) Simplicity is preserved under Morita equivalence

(3) Let $S = \mathrm{End}_D(V)$. Then $T = \mathrm{End}_{R\otimes D^\circ}(V)$, since V is a faithful S-module. Now, $R\otimes D$ is a simple ring, so $R\otimes D = \mathrm{End}_E(W)$ for a unique simple module W and $E = \mathrm{End}_{R\otimes D}(W)$. In particular, $V = W^m$ for some m. Hence:

$$[T : k] = [\mathrm{End}_{R\otimes D}(V) : k] = [\mathrm{End}_{R\otimes D}(W^m) : k] = \\ = [M_m(E) : k] = m^2[E : k]$$

On the other hand,

$$[R : k][D : k] = [R\otimes D : k] = [W : E]^2[E : k] = \frac{[W : k]^2}{[E : k]}$$

and

$$[S : k] = [D : k][V : D]^2 = [D : k](\frac{[V : k]}{[D : k]})^2 = \\ = \frac{[W^m : k]^2}{[D : k]} = \frac{m^2[W : k]^2}{[D : k]}$$

The result now follows from arithmetic. □

Corollary 6.2.8. *If S and $R \subseteq S$ are CSAs, then the map $R\otimes Z_S(R) \to S$ given by $r\otimes t \mapsto rt$ is an isomorphism*

Corollary 6.2.9 (The double centraliser theorem)**.** *For any simple subalgebra $R \subseteq S$ of a CSA S, $Z_S(Z_S(R)) = R$*

Corollary 6.2.10. *Let D be a central division algebra of dimension n^2 over k. Then any maximal subfield of D has dimension n over k*

Proof. Let $L \subseteq D$ be a maximal subfield. If $a \in Z_D(L)$, then $L[a]$ is also a subfield of D, hence $a \in L$. So L is equal to its centraliser, and by the dimension statement its dimension is n. □

6.2.11. *Finite division rings.* It turns out that they are all trivial:

Theorem 6.2.12 (Wedderburn's little theorem)**.** *If k is a finite field, then* $\mathrm{Br}(k) = 0$*. In other words, every finite division ring is a field*

Proof. Let D be a central division algebra over k, and for every $x \in D$, let $k \subset L_x \subset D$ be a maximal subfield containing x. We know that the dimension of each L_x is the same, hence, since k is finite, they are all isomorphic. By the Skolem–Noether Theorem, they are all conjugate. In particular, the multiplicative group D^* is a union of conjugates of ${L_0}^*$. The following lemma implies that $L_0 = D$. □

Lemma 6.2.13. *If G is a finite group and H is a proper subgroup, then G is not a union of conjugates of H.*

End of lecture 22, Jun 21

Proof. Consider the action of G on subgroups by conjugation. The stabiliser $N(H)$ of H contains H, and the conjugates are given by $G/N(H)$. Hence the number of elements in the union of all conjugates is at most $m = |G/N(H)||H|$, and is strictly smaller if H is a proper subgroup, since all conjugates have at least 1 in common. On the other hand, the number of elements in G is $|G/H||H|$ so is at least m. □

6.2.14. *Affine geometry.* We describe an interesting relation between the division rings and classical geometry. One way to make Euclid's axioms more precise is to work with structure consisting of a set P of "points", a set L of "lines", and a relation $I \subseteq P \times L$ called "incidence". We say that $p \in P$ belongs to $l \in L$ (or that l contains p) if $I(p, l)$ holds. We say that lines $l_1, l_2 \in L$ are *parallel* if $l_1 = l_2$ or no point belongs to both.

An *affine plane* is a structure of this form that satisfies the following axioms:

affine plane

(1) For any $x \neq y \in P$, there is a unique $l \in L$ containing both. We denote this line by $l(x, y)$.
(2) For any $l \in L$ there are $x \neq y \in P$ belonging to l
(3) There are three points not on belonging to the same line.
(4) For any $p \in P$ and any $l \in L$ there is a unique $l' \in L$ containing p and parallel to l (this is the *parallel axiom*)

A *coordinate system* on an affine plane $\mathbb{A}$ is a choice of ordered triple $(0, e_1, e_2)$ of non-colinear points, as in the third axiom. A *coordinatised affine plane* is an affine plane along with a coordinate system.

An example of an affine plane is the usual Euclidean plane $P = \mathbb{R}^2$, with lines given by linear equations. We view it as coordinatised by the triple $((0, 0), (1, 0), (0, 1))$. In fact, the same choice produces from each division ring D a coordinatised affine plane $\mathbb{A}(D)$. It is not the case that every affine plane is of this form, since all such planes satisfy an additional property, called *Desargues's Theorem.* However, we have:

Desargues's Theorem

Fact 6.2.15. *The assignment $D \mapsto \mathbb{A}(D)$ forms an equivalence of categories between division rings and coordinatised affine planes satisfying Desargues's Theorem:*

> *Assume we are given distinct points a, a', b, b', c, c' such that $l(a, a') \neq l(b, b'), l(c, c')$ meet at the same point, and such that $l(a, b), l(a', b')$ are parallel and $l(a, c), l(a', c')$ are parallel. Then $l(b, c)$ and $l(b', c')$ are parallel as well.*

Hence, the theory of division algebras that we are studying is, in some sense, the theory of Desarguesian affine planes. For example, it immediately follows that in a finite such plane the number of points is a square (though this is easy to prove in an arbitrary affine plane). Given this result, it is reasonable to ask which affine planes correspond to fields. The answer is given by another classical result, *Pappus's Theorem*:

Pappus's Theorem

Fact 6.2.16. *The assignment $F \mapsto \mathbb{A}(F)$ forms an equivalence of categories between fields and coordinatised affine planes satisfying Pappus's Theorem:*

> *Assume we are given distinct points a, a', b, b', c, c' such that a, b, c are on the same line, a', b', c' are on the same line, $l(a, b')$ is parallel to $l(a', b)$ and $l(a, c')$ is parallel to $l(a', c)$. Then $l(b, c')$ is parallel to $l(b', c)$.*

These two facts show immediately that the Pappus property implies the Desargues property, which is not immediately obvious even in the usual real plane, but also, using Theorem 6.2.12, we obtain:

Corollary 6.2.17. *Any finite Desarguesian affine plane satisfies that Pappus property.*

There is no known direct proof of this (without passing to division rings). See [Har00] or [Bal+07] for more details on this.

6.3. **Descent theory.** We now take more seriously the point of view that a CSA is "an object that becomes trivial after base change". The first step is to pass to finite Galois extensions.

6.3.1. *Splitting fields.* We first note that, while all CSAs become trivial over an algebraic closure of k, each CSA becomes trivial over a finite extension: If $f : k^a \otimes A \to M_n(k^a)$ is an isomorphism, corresponding to a map $f_0 : A \to M_n(k^a)$, the f_0 will determine an isomorphism over any field that contains the image of a basis of A over k.

splitting field
$\mathrm{Br}(L/k)$

A field L such that $L \otimes A$ is trivial is called a *splitting field* for A. We denote by $\mathrm{Br}(L/k)$ the subgroup of $\mathrm{Br}(k)$ consisting of elements represented by algebras that split over L.

Since each CSA is an endomorphism algebra over a central division algebra, splitting CSAs is equivalent to splitting division algebras. By taking the normal closure of a finite splitting field, we see that each CSA has a finite normal splitting field. More interestingly, we may take the splitting field to be separable (hence Galois). To prove that, we first note:

Proposition 6.3.2. *If $k \subset L \subset A$ is a maximal subfield of a central simple algebra A, then L splits A*

Proof. We have an algebra homomorphism $f : L \otimes A^\circ \to \mathrm{End}_L(A)$ given by $f(x \otimes a)(b) = xba$. Since A is simple, so is $L \otimes A^\circ$, so f is injective. By Corollary 6.2.10, the dimension of both sides is the same, so f is an isomorphism. □

As mentioned above, splittings CSAs amounts to splitting division algebras. For them, we have:

Proposition 6.3.3. *Let D be a central division algebra over k.*

(1) *(**The Jacobson–Noether Theorem**) If $D \neq k$, then D contains an element separable over k (but not in k)*

*(2) (**Koethe's Theorem**) Any separable extension $k \subseteq L \subset D$ is contained in a maximal extension which is separable.*

Proof. (1) Assume not. Then there is some $a \in D \setminus k$ with $a^p \in k$ (where p is the characteristic of k). Let $d : D \to D$ be the k-linear map $d(x) = xa - ax$, and let $C = Z_D(()a)$. We note that d is a derivation over C: $d(xy) = d(x)y + xd(y)$ for all $x, y \in D$, and $d(x) = 0$ for $x \in C$ (so that d is C-linear). One sees by induction that $d^n(x) = \sum_{i=0}^{n}(-1)^i\binom{n}{i}a^i x a^{n-i}$. In particular, $d^p(x) = xa^p - a^p x = 0$ (since $a^p \in k$).

Since $d \neq 0$, there is some $w \in D \backslash C$. By applying d enough times, we may assume $d^2(w) = 0$, and by dividing by $d(w) \in C$, we may assume $d(w) = 1$. Hence $wa - aw = 1$, or $c := wa = 1 + aw = 1 + a(wa)a^{-1} = 1 + aca^{-1}$. Consider the subfield $k(c)$. By assumption, this is a purely inseparable extension, so $c^{p^e} \in k$ for some e. Applying this to the equation, we get $c^{p^e} = 1 + ac^{p^e}a^{-1} = 1 + c^{p^e}$ since k is the centre. This is a contradiction.

(2) Let $D_0 = Z_D(L)$. By Corollary 6.2.9, $Z(D_0) = L$. Now apply the previous part to D_0. □

Corollary 6.3.4. *Every CSA has a finite Galois splitting field. Hence,* $\mathrm{Br}(k) = \cup_L \mathrm{Br}(L/k)$*, where the union ranges over finite Galois extensions.*

The corollary allows us to concentrate on $\mathrm{Br}(L/k)$, and therefore to use Galois theory. **From now till the end of the section, we fix a finite Galois extension L of k, and denote by $\Gamma = \mathrm{Aut}(L/k)$ the Galois group.**

6.3.5. *Forms and torsors.* We are interested in the following situation: We have a category of objects defined over a field k, which we can "base change" to a field extension L. We would like to classify such objects Y over k that becomes isomorphic to a particular object X after base change to L. Such a Y is called a *form* of X. Let P be the set of isomorphisms from the base change Y_L to the base change X_L. This set is non-empty precisely if Y_L is, in fact, isomorphic to X_L. We assume the base change to be functorial in L, so that automorphism group Γ of L over k acts on X_L and on Y_L. This action is not by L-automorphisms, since it is not L-linear, but each isomorphism $f : Y_L \to X_L$ is L-linear, so if $\gamma \in \Gamma$, the composed map $\gamma_X \circ f \circ {\gamma_Y}^{-1}$ is again in P.

form

We have been vague about our notion of "base change", as well as about linearity, etc. It is possible to make such an abstract theory precise, but we will limit ourselves to two instances of it below. At any rate, we obtain a set P along with an action of the group Γ on it. Taking $Y = X$, the set of automorphisms of X_L is a group under composition, and Γ acts on it by group automorphisms. Furthermore, G acts on P via composition, and this action again commutes with the action of Γ on G and on P. We arrive at the following definition:

torsor

Definition 6.3.6. Let Γ be a group, and let G be a group on which Γ acts by automorphisms. A *torsor* over G is a non-empty set P with an action of Γ on P along with an action $m : G \times P \to P$ commuting with Γ, such that the combined map $(m, \pi) : G \times P \to P \times P$ (with π the second projection) an isomorphism.

As usual, abbreviate the action $m(g, p)$ as gp. The condition thus says that $(g, p) \mapsto (gp, p)$ is a bijection. In other words, for any two element $p, q \in P$, there is a unique element $g \in G$ with $gp = q$. In this definition there are no fields, and the fact that Γ is a Galois group is irrelevant. In fact, the same definition can be made in any category with finite products. But we will apply it in the case that Γ is a Galois group, as outlined above.

trivial torsor

$H^1(\Gamma, G)$

A map of torsors $t : P \to Q$ is a map of sets that commutes with the actions of G and Γ. Such a map is necessarily an isomorphism. The *trivial torsor* is $P = G$, with the action given by the group operation. A torsor is isomorphic to the trivial torsor if and only if it has a point fixed by Γ: $1 \in G$ is such a point, and if $p \in P$ is fixed, there is a unique isomorphism $t : G \to P$ determined by $t(1) = p$. Thus, if Γ is trivial, all torsors are isomorphic, but in general, they need not be. We denote the set of isomorphism classes of G-torsors by $H^1(\Gamma, G)$. It is a *pointed set*, i.e., it has a special element $*$, the class of the trivial torsor.

As outlined above, a form of X gives rise to a torsor P over the group G of automorphisms of X_L, with the action of the Galois group Γ, the set of isomorphisms from Y_L to X_L. This torsor is trivial precisely if Y was already isomorphic to X over k. Conversely, under some assumptions on the category, a torsor over G gives rise to a form of X. For example, if all colimits exist, we may construct such a form as $\varinjlim_{a \in P} X_a$, where X_a is a copy of X_L, and for $a, b \in P$, the map from X_a to X_b is given by the action of g on X_L, where $ga = b$. We will see a concrete version of this below.

End of lecture 23, Jun 22

semi-linear action

$H^0(\Gamma, V_L)$

6.3.7. *Forms of a vector space.* One of the main examples of the general theory outlined above is for finite dimensional vector spaces. Let V be a finite dimensional vector space over a field k. Then Γ acts on the base change $V_L = L \otimes V$ by acting on L, but this is not an L-linear action. Instead, we have a *semi-linear action*, with $\gamma(av) = \gamma(a)\gamma(v)$ for all $\gamma \in \Gamma$, $a \in L$ and $v \in V_L$. If we are given V_L with this action, we may recover the space V we started with by taking the invariants, which we denote by $H^0(\Gamma, V_L)$.

twisted group algebra

L_Γ

We are interested in the converse: Given a space W over L, and a semi-linear action of Γ on W, is it the case that it comes from a vector space over k? In considering linear actions of a group (i.e., representations), it was fruitful to consider the group algebra. Similarly, here we consider the *twisted group algebra* L_Γ, which is the L-vector space spanned by Γ over L, with product determined by $\gamma \cdot a = \gamma(a)\gamma$. It is easy to see that a semi-linear action of Γ on W is the same as an L_Γ-module structure, over L.

Proposition 6.3.8. *The ring L_Γ is simple*

Proof. Assume that $t : L_\Gamma \to R$ is a ring homomorphism. Since both the elements of Γ and the non-zero elements of L are invertible, none of the maps to 0, and we may view L as a subfield of R.

Let $\sum a_\gamma \gamma$ be an element in the kernel, with minimum number of summands. By the remark above, this number is at least 2. By multiplying with some γ on the right and $\frac{1}{a_\gamma}$ on the left, we may assume $a_e = 1$. Multiplying by some $b \in L^*$ on the right, we obtain that $\sum a_\gamma \gamma(b)\gamma$ goes to 0 as well. By minimality, it follows that $\gamma(b)$ is independent of γ, hence $\gamma(b) = b$ for all γ that appear in this sum (since e appears in it). Hence, the subgroup generated by these γ fixes all elements of L. This contradicts the fact that L is Galois. □

Corollary 6.3.9. *Let L be a finite Galois extension of k with Galois group Γ. The base change functor $V \mapsto L \otimes V$ is an equivalence of categories from $\mathcal{V}ec_k$ to the category $\Gamma - \mathcal{V}ec_L$ of finite dimensional L-vector spaces with semi-linear Γ action. An inverse is given by $W \mapsto H^0(\Gamma, W)$*

Proof. View $\mathcal{C} = \Gamma - \mathcal{V}ec_L$ as the category of finitely-generated L_Γ-modules. Since L_Γ is simple (and Artinian), any non-zero module is a projective generator. Taking L to be such, Morita theory shows that the functor $W \mapsto \mathrm{Hom}_{L_\Gamma}(L, W)$ is an equivalence with the category of finitely-generated modules over $R = \mathrm{End}_{L_\Gamma}(L)^\circ$. An element f of $\mathrm{Hom}_{L_\Gamma}(L, W)$ is determined by $f(1) \in W$, and since 1 is fixed by Γ, $w \in W$ determines such a map if and only if it is fixed by Γ. Thus, the functor is given by $W \mapsto H^0(\Gamma, W)$, and in particular, $R = H^0(\Gamma, L) = k$. □

Corollary 6.3.10 (Hilbert's Theorem 90)**.** *For any n, $H^1(\Gamma, \mathrm{GL}_n(L)) = 0$*

Proof. Let P be a torsor over $G = \mathrm{GL}_n(L)$, with a compatible action of Γ. Let $W = \{f : P \to L^n \mid f(gp) = gf(p)\}$. Then W is an L-vector space, and if $p \in P$ is any point, $f \mapsto f(p)$ is an isomorphism of W with L^n, so it has dimension n. The rule $\gamma(f)(p) = \gamma(f(\gamma^{-1}(p)))$ defines a semi-linear Γ action on W, so by the above, there are n k-independent elements of W commuting with the action of Γ, $f_1, \dots, f_n$. Then $f = (f_1, \dots, f_n)$ is a Γ-isomorphism from P to the trivial torsor G. □

End of lecture 24, Jun 28

Remark 6.3.11. One unfortunate aspect of the theory just described is that it works only for Galois extensions. For example, if L is a purely inseparable extension, then Γ is trivial, and so we cannot hope to recover a vector space over k from a vector space over L via a Γ-action.

To overcome this deficiency, consider again the case where L is Galois over k. A k-linear action of Γ on an L-vector space W can be described via a k-linear map $W \to W^\Gamma$ (W^Γ is the space of maps of sets), sending w to $\gamma \mapsto \gamma(w)$. We have L-linear isomorphisms $W^\Gamma \to W \otimes_L L^\Gamma \to W \otimes_L L \otimes L \to W \otimes L$, where the middle one is by Prop. 3.6.6. The action being semi-linear corresponds to the composed map $W \to W \otimes L$ being linear with respect to the L-structure on the second component L of $W \otimes L$. In other words, it is a map $d : L \otimes W \to W \otimes L$ over $L \otimes L$. The fact that the map is coming from a group action is equivalent to the conditions that d is an isomorphism, that it commutes with the vector space multiplication $L \otimes W \to W$ on W, and that $(d \otimes \mathbf{1}) \circ (\mathbf{1} \otimes d) : L \otimes L \otimes W \to W \otimes L \otimes L$ is the map given by d on the first and third component (and the identity on the middle one).

descent datum

A *descent datum* on W is a map d as above. The previous paragraph explains that when L is Galois over k, giving a such a datum is equivalent to giving a semilinear action of Γ on W. Hilbert 90 then says that the category of L spaces with descent data (and obvious morphisms) is equivalent to the category of k-vector spaces. However, to formulate the definition of descent datum, and to state the claim, the assumption that L is Galois is no longer needed. Indeed, the statement turns out to be true, and is part of much more general theory known as *faithfully flat descent* (which is, itself, a special case of *Beck's theorem*, a purely categorical statement). □

faithfully flat descent

6.3.12. *Descent for CSAs.* Suppose that $A = M_n(k)$ is a matrix algebra over k. Then $B = L \otimes A = M_n(L)$ is again endowed with an action of Γ, which is a semi-linear action by ring automorphisms, and $H^0(\Gamma, B) = A$. Assume, as before, that we are given a matrix algebra B over L with a semilinear Γ-action. Is $H^0(\Gamma, B)$ necessarily a matrix algebra over k? The fact that we have a non-trivial theory of CSAs means that the answer is no. For example, the base change of the Quaternions $\mathbb{H}$ to $\mathbb{C}$ is a matrix algebra.

However, we have:

Proposition 6.3.13. *Let L be a finite Galois extension of k with Galois group Γ. The base change functor $A \mapsto L \otimes A$ is an equivalence of categories from CSAs over k to CSAs over L along with a semilinear Galois action, with inverse $B \mapsto H^0(\Gamma, B)$.*

Proof. We know by Hilbert 90 that the map $L \otimes H^0(\Gamma, B) \to B$ is an isomorphism of vector spaces. Since L is the centre of B, it commutes with the image of $H^0(\Gamma, B)$, hence it is also an isomorphism of algebras. If $I \subset H^0(\Gamma, B)$ is a two-sided ideal, then (again, since L is central) $L \otimes I$ is an ideal in B, hence is 0. Since k is a field, this implies that $I = 0$. Finally, we have $L = Z(B) = Z(L \otimes H^0(\Gamma, B)) = L \otimes Z(H^0(\Gamma, B))$ by Prop. 6.1.2. Hence $Z(H^0(\Gamma, B)) = k$. □

To reformulate the result in terms of H^1, we need to identify the automorphisms of the trivial algebra $A = M_n(L)$. By the Skolem–Noether theorem (6.2.4), every automorphism is given by conjugation by an invertible matrix. Thus, we have a surjective map from $\mathrm{GL}_n(L)$ to the group of automorphisms of A. The kernel is, by definition, the centre of $\mathrm{GL}_n(L)$, which is the group of scalar matrices. Thus, the automorphism group is the *projective general linear group* $\mathrm{PGL}_n(L) = \mathrm{GL}_n(L)/\,\mathrm{G_m}(L)$ (where $\mathrm{G_m}(L) = \mathrm{GL}_1(L)$ is the multiplicative group of L). Hence:

projective general linear group

Corollary 6.3.14. *The set $H^1(\Gamma, \mathrm{PGL}_n(L))$ is canonically identified with the set of isomorphism classes of n^2-dimensional CSAs over k*

Example 6.3.15. Let $k = \mathbb{R}$, $L = \mathbb{C}$ and $A = \mathbb{H}$. Then A defines the element of $H^1(\Gamma, \mathrm{PGL}_2(\mathbb{C}))$ in which complex conjugation $c \in \Gamma$ acts as $c(\left(\begin{smallmatrix} x & y \\ z & w \end{smallmatrix}\right)) = \left(\begin{smallmatrix} -w & y \\ z & -x \end{smallmatrix}\right)$ (this follows from Exercise 6.2.2). The theorem of Frobenius implies that this is the only non-trivial element of $H^1(\Gamma, \mathrm{PGL}_2(\mathbb{C}))$. □

6.3.16. *Exact sequences.* The set $H^1(\Gamma, G)$ has a special point $*$, corresponding to the trivial torsor G. In other words, it is a pointed set. We say that a sequence of maps $\dots \xrightarrow{t_i} A_i \xrightarrow{t_{i+1}} A_{i+1} \to \dots$ of pointed sets is *exact* if for all i, the image of t_i is the fibre (inverse image) of $*$ under t_{i+1}. **exact**

The assignment $G \mapsto H^1(\Gamma, G)$ is a functor into the category of pointed sets, as follows: If $f : G \to H$ is a group homomorphism (commuting with the action of Γ), we map the G torsor P to the H torsor $H \otimes_G P = H \times P / \sim$, with $(h, gp) \sim (hf(g), p)$ for all $h \in H$, $g \in G$ and $p \in P$, and with Γ acting on both components.

Proposition 6.3.17. *Let H be a quotient group of G, with kernel K (all compatible with the action of Γ). Then the resulting sequence $H^1(\Gamma.K) \to H^1(\Gamma, G) \to H^1(\Gamma, H)$ is exact.*

Proof. Let P be a K-torsor. Then $H \otimes_G (G \otimes_K P) = H \otimes_K P$. The equivalence relation on the last torsor identifies (h, p) with (h, q) for all $p, q \in P$ (since K goes to $e \in H$), so in other words, with H.

Assume that P is a G-torsor. Then $H \otimes_G P$ is the set P/K of K-orbits on P. If this is trivial, it means that there is such an orbit P_0 that is fixed by Γ. It is easy to see that the action map $G \otimes_K P_0 \to P$ is an isomorphism. □

We would like extend the above sequence on both sides. First, for any Γ-group G, recall that $H^0(\Gamma, G)$ is the set of fixed points for the action of Γ. This is actually a group, but for now we will ignore the group structure, and view it as a pointed set, with the identity as the special point.

Proposition 6.3.18. *If $\pi : G \to H$ is surjective with kernel K, there is a long exact sequence*

$$\begin{aligned} 0 \to H^0(\Gamma, K) \to & H^0(\Gamma, G) \to H^0(\Gamma, H) \xrightarrow{d} \\ & H^1(\Gamma, K) \to H^1(\Gamma, G) \to H^1(\Gamma, H) \end{aligned}$$

Proof. The map d is defined as follows: if $h \in H$ is fixed by Γ, then the fibre $\pi^{-1}(h)$ is a coset of K in G, fixed (as a set) by Γ. Hence, it is a K-torsor, and we define $d(h)$ to be its class. This torsor is trivial precisely if this fibre has a Γ-fixed point, i.e., if h is in the image of $H^0(\Gamma, G)$.

For h as above, the G-torsor $G \otimes{}_K \pi^{-1}(h)$ is trivial, as witnessed by the class of (g^{-1}, g), for any $g \in \pi^{-1}(h)$. Conversely, if P is a K-torsor with $G \otimes_K P$ trivial, witnessed by Γ fixing the class of (g, p), then P is isomorphic to the K-coset of g^{-1} inside G. The verification and other claims are left as an exercise. □

Example 6.3.19. Consider the case $G = \mathrm{GL}_n(L)$, $H = \mathrm{PGL}_n(L)$ and $K = \mathrm{G_m}(L)$. Hilbert 90 asserts that the codomain of d is trivial in this case, so every element in $\mathrm{PGL}_n(L)$ fixed by Γ is represented by some element in $\mathrm{GL}_n(k)$. For example, if $k = \mathbb{R}$ and $L = \mathbb{C}$, the class of the matrix $\left(\begin{smallmatrix} i & 0 \\ 0 & -i \end{smallmatrix}\right)$ is fixed by complex conjugation, so is represented by a matrix with real entries. □

The notation suggests that the sequence above can be continued with entries involving H^2. This will be true once we define these terms, but the kind of geometric definition we used to define H^1 (via torsors) becomes more complicated in this case, and we will not give it. Instead, we concentrate on the case where the groups are abelian. In this case, we will give a more computational definition, which works for H^i in general.

6.3.20. *Cocycles and coboundaries.* Let P be a torsor over G, and let $p \in P$ be an arbitrary point. For every $\gamma \in \Gamma$, there is a unique $g_\gamma \in G$ with $\gamma(p) = g_\gamma p$ (since P is a torsor). We have $g_1 = e$, and if $\tau \in \Gamma$ as well, we have

$$g_{\tau\gamma} p = \tau(\gamma(p)) = \tau(g_\gamma p) = \tau(g_\gamma)\tau(p) = \tau(g_\gamma) g_\tau p$$

Again, since P is a torsor, we may "cancel" p, so $g_{\tau\gamma} = \tau(g_\gamma) g_\tau$. By setting $f(\gamma) = {g_\gamma}^{-1}$, we thus obtain a function $f : \Gamma \to G$ satisfying

$$f(\tau\gamma) = f(\tau)\tau(f(\sigma)) \tag{6.1}$$

The function f depends on the choice of p: if $p' \in P$ is another point, there is a (unique) element $h \in G$ such that $p' = hp$. If f' corresponds to p' we have:

$$hp = p' = f'(\gamma)\gamma(p') = f'(\gamma)\gamma(hp) = f'(\gamma)\gamma(h)\gamma(p) = f'(\gamma)\gamma(h)f(\gamma)^{-1}p$$

Once again, we are allowed to cancel p and obtain

$$f'(\gamma) = hf(\gamma)\gamma(h)^{-1} \tag{6.2}$$

$Z^1(\Gamma, G)$

Thus, if we let $Z^1(\Gamma, G)$ be the set of functions $f : \Gamma \to G$ satisfying (6.1), and declare $f \sim f'$ if there is some $h \in G$ such that (6.2) is satisfied for all γ, we find that each torsor P determines an element f_P of $Z^1(\Gamma, G)/\sim$. It is easy to see that isomorphic torsors determine the same class, so we proved one direction of:

Proposition 6.3.21. *The assignment $f \mapsto f_P$ determines a bijection of pointed sets $H^1(\Gamma, G) \to Z^1(\Gamma, G)/\sim$*

Proof. We construct a map back. Let $f \in Z^1(\Gamma, G)$. We let $P = G$ as a set with a G-action, and we define the action of Γ by the same formula: $\gamma_P(g) = \gamma_G(g)f(\gamma)^{-1}$. The verification that this is well defined and forms a two-sided inverse is left as an exercise. □

The trivial torsor can be represented by the constant function with value $e \in G$. The equivalence class of this function consists of the functions of the form $f_h : \gamma \mapsto \gamma(h)h^{-1}$ for some fixed h. The subset of all such functions is denoted by $B^1(\Gamma, G)$.

$B^1(\Gamma, G)$

Assume now that G is abelian. In this case, the condition (6.1) is preserved under pointwise multiplication and inverse, so $Z^1(\Gamma, G)$ is a subgroup of the group of all functions from Γ to G. Furthermore, $f \sim f'$ if and only if the differ by a function in $B^1(\Gamma, G)$, which forms a subgroup of $Z^1(\Gamma, G)$. We thus obtain:

Corollary 6.3.22. *If G is an abelian group with a Γ-action, then $H^1(\Gamma, G) = Z^1(\Gamma, G)/B^1(\Gamma, G)$*

We note that so far, we had no group structure on H^1, this is a special feature of the case where G is abelian. In this case, it not hard to describe the group structure in terms of torsors as well:

Exercise 6.3.23. Assume that G is abelian. For G torsors P and Q, define $P \otimes_G Q$ to be the quotient $P \times Q/\sim$, where $(gp, q) \sim (p, gq)$ for all $g \in G$, with G acting on either component, and Γ acting on both. Define P^{-1} to be the torsor with the same Γ action, but with $g \in G$ acting as g^{-1}. Show that this determines a well defined group structure on $H^1(\Gamma, G)$, which makes the map above a group isomorphism.

We now note the following: The subgroup $B^1(\Gamma, G)$ is the image of the map $h \mapsto f_h$. Its kernel is the set $Z^0(\Gamma, G) = H^0(\Gamma, G)$, the elements of G fixed by Γ. On the other hand, $Z^1(\Gamma, G)$ is the kernel of the map $d^1 : G^\Gamma \to G^{\Gamma\times\Gamma}$ given by $d^1(f)(\tau, \gamma) = \tau(f(\gamma))f(\tau\gamma)^{-1}f(\tau)$. The general definition generalises this observation as follows:

Definition 6.3.24. Let Γ be a group acting on an abelian group G, and let $n \geq 0$

(1) The group of *n-cochains*, $C^n(\Gamma, G)$, is the set G^{Γ^n}, with pointwise multiplication. We let $C^{-1} = 0$, the trivial group.

n-cochains
$C^n(\Gamma, G)$

(2) The *coboundary homomorphism* $d^n : C^{n-1} \to C^n$ is given by

coboundary homomorphism

$$d^n(f)(\gamma_1, \dots, \gamma_n) = \gamma_1(f(\gamma_2, \dots, \gamma_n)) \prod_{i=1}^{n} f(\delta_i(\bar{\gamma}))^{(-1)^i}$$

where $\delta_i(\bar{\gamma}) = (\gamma_1, \dots, \gamma_i\gamma_{i+1}, \dots, \gamma_n)$ for $i < n$, and δ_n is the projection to the first $n-1$ arguments.

(3) The *n-cocycles* $Z^n(\Gamma, G)$ are the kernel of d^{n+1}, and the *n-coboundaries* $B^n(\Gamma, G)$ are the image of d^n. The n-th *cohomoology group* $H^n(\Gamma, G)$ is the quotient group $Z^n(\Gamma, G)/B^n(\Gamma, G)$

n-cocycles
$Z^n(\Gamma, G)$
n-coboundaries
$B^n(\Gamma, G)$
cohomoology group
$H^n(\Gamma, G)$

Remark 6.3.25. (1) It is easy to check that d^n is indeed a homomorphism, and that $d^{n+1} \circ d^n = 0$, so the definition makes sense.

(2) We saw above that the terminology is consistent with our previous definition in the case $n = 0, 1$.

□

End of lecture 25, Jun 29

This definition can be viewed as a special case of the *derived functor formalism* (for the functor of invariants), which provides an indefinitely long

derived functor formalism

exact sequence in the style of Prop. 6.3.18 for any short exact sequence of *abelian* groups. We will not go into that, but instead prolong the sequence by one more step, in the case the kernel is central:

Proposition 6.3.26. *Assume that* $\pi : G \to H$ *is a surjective homomorphism of* Γ*-groups, whose kernel* K *is in the centre of* G*. Then there is a map* $c : H^1(\Gamma, H) \to H^2(\Gamma, K)$ *making the sequence*

$$H^1(\Gamma, G) \to H^1(\Gamma, H) \xrightarrow{c} H^2(\Gamma, K)$$

of pointed sets exact.

Proof. Let P be an H-torsor. Then G acts on P transitively, with stablilisers K. If $p \in P$ and $\gamma \in \Gamma$, there is an element $g_\gamma \in G$ (not unique), such that $\gamma(p) = g_\gamma p$. If $\tau \in \Gamma$ is another element, it need not be the case that $g_{\tau\gamma} = \tau(g_\gamma)g_\tau$, but this is true after applying π (since we obtain a cocycle corresponding to P). Hence, the function $f : \Gamma \times \Gamma \to G$ given by $f(\tau, \gamma) = g_{\tau\gamma}\tau(g_\gamma)^{-1}{g_\tau}^{-1}$ takes values in K, so determines an element of $C^2(\Gamma, K)$. If f' is obtained by choosing different representatives g'_γ, the ratio $f^{-1}f'$ is a coboundary. The verification that the map is well defined, and takes values in the cocycles is left as an exercise. If P comes from a torsor Q over G, we may choose f above to be an actual cocycle for Q, which will therfore go to 0 in $H^2(\Gamma, K)$, and conversely, if $c(f) = 0$, then f is a cocycle. □

For each $n > 0$, let $H^2_n(\Gamma, \mathrm{G_m}(L))$ be the image of the map

$$c_n : H^1(\Gamma, \mathrm{PGL}_n(L)) \to H^2(\Gamma, \mathrm{G_m}(L))$$

provided by the proposition for the case $H = \mathrm{PGL}_n(L)$, $G = \mathrm{GL}_n(L)$ and $K = \mathrm{G_m}(L)$, and let $\bar{H}^2(\Gamma, \mathrm{G_m}(L)) = \bigcup_n H^2_n(\Gamma, \mathrm{G_m}(L))$. Then we have:

Corollary 6.3.27. $\bar{H}^2(\Gamma, \mathrm{G_m}(L))$ *is a subgroup of* $H^2(\Gamma, \mathrm{G_m}(L))$*, isomorphic to* $\mathrm{Br}(L/k)$

Proof. Let A and B be two CSAs over k, of dimensions n^2 and m^2. If f_A and f_B are cocycles representing the classes of A and B in $H^1(\Gamma, \mathrm{PGL}_n)$ and $H^1(\Gamma, \mathrm{PGL}_m)$, respectively, then a direct computation shows that $f_A \otimes f_B$ represents $A \otimes B$, where $(f_A \otimes f_B)(\gamma) = f_A(\gamma) \otimes f_B(\gamma) \in \mathrm{Aut}(A \otimes B)$. Using the explicit description of c_n from the proposition, we get that $c_{nm}(f_A \otimes f_B) = c_n(f_A)c_m(f_B)$. Hence, on the level of isomorphism classes, we obtain a map $t([A]) = c_n(f_A)$, with $t([A \otimes B]) = t([A])t([B])$. This map is injective Prop. 6.3.26, since $H^1(\Gamma, \mathrm{GL}_n) = *$ by Hilbert 90.

Now, if A and B are Morita equivalent (i.e., give the same element in $\mathrm{Br}(L/k)$), then for some matrix algebras A' and B', $A \otimes A'$ and $B \otimes B'$ are isomorphic (by Prop. 6.1.11). Hence $t(A) = t(A \otimes A') = t(B \otimes B') = t(B)$. It follows that t induces an injective group homomorphism $t : \mathrm{Br}(L/k) \to H^2(\Gamma, \mathrm{G_m}(L))$. By definition, the image is $\bar{H}^2(\Gamma, \mathrm{G_m}(L))$. □

We may apply this result as follows:

Proposition 6.3.28. *The image of the map*

$$c_n : H^1(\Gamma, \mathrm{PGL}_n(L)) \to H^2(\Gamma, \mathrm{G_m}(L))$$

is killed by n. In particular, the Brauer group $\mathrm{Br}(L/k)$ is torsion.

Sketch of proof. We assume for the proof that n is prime to the characteristic. Let $\det : \mathrm{GL}_n \to \mathrm{GL}_n$ be the determinant map. Its kernel is (by definition) SL_n, and its restriction to the centre $\mathrm{G_m}$ is raising to power n. By extending L to be closed under n-th roots (since n is prime to p, this is a separable extension), make the columns in the following diagram exact:

$$\begin{tikzcd}
 & 1 \arrow[d] & 1 \arrow[d] & & \\
1 \arrow[r] & \mu_n \arrow[r] \arrow[d] & \mathrm{SL}_n \arrow[r] \arrow[d] & \mathrm{PGL}_n \arrow[r] \arrow[d, "="] & 1 \\
1 \arrow[r] & \mathrm{G_m} \arrow[r] \arrow[d, "x \mapsto x^n"] & \mathrm{GL}_n \arrow[r] \arrow[d, "\det"] & \mathrm{PGL}_n \arrow[r] & 1 \\
 & \mathrm{G_m} \arrow[r, "="] \arrow[d] & \mathrm{G_m} \arrow[d] & & \\
 & 1 & 1 & &
\end{tikzcd}$$

Here, μ_n is the group of n-th roots of 1. It follows that the quotient SL_n / μ_n is indeed PGL_n. Comparing the long exact sequences for the rows, we have a commutative diagram

$$\begin{tikzcd}
H^1(\Gamma, \mathrm{PGL}_n) \arrow[r] \arrow[d, "="] & H^2(\Gamma, \mu_n) \arrow[d] \\
H^1(\Gamma, \mathrm{PGL}_n) \arrow[r, "c_n"] & H^2(\Gamma, \mathrm{G_m})
\end{tikzcd}$$

So c_n factors through $H^2(\Gamma, \mu_n)$. A direct calculation shows that this group is n-torsion. □

Remark 6.3.29. The proof above works for general n if we use faithfully flat descent in place of Galois theory, and a suitable notion of μ_n.In fact, it is known (and not hard to prove) that $\mathrm{Br}(L/k) = H^2(\Gamma, \mathrm{G_m})$. See [Mil80, §IV.2] for details (and a much more general theory). □

References

[Art] Michael Artin. *Noncommutative rings.* URL: `http://math.mit.edu/~etingof/artinnotes.pdf`.

[Bal+07] Philippe Balbiani et al. "Logical theories for fragments of elementary geometry". In: *Handbook of spatial logics.* Springer, Dordrecht, 2007, pp. 343–428. DOI: `10.1007/978-1-4020-5587-4_7`.

[Bas68] Hyman Bass. *Algebraic K-theory.* W. A. Benjamin, Inc., New York-Amsterdam, 1968, pp. xx+762.

[Cla] Pete Clark. *Noncommutative algebra.* URL: `http://math.uga.edu/~pete/noncommutativealgebra.pdf`.

[FD93] Benson Farb and R. Keith Dennis. *Noncommutative algebra.* Vol. 144. Graduate Texts in Mathematics. Springer-Verlag, New York, 1993, pp. xiv+223. ISBN: 0-387-94057-X. DOI: `10.1007/978-1-4612-0889-1`.

[Har00] Robin Hartshorne. *Geometry: Euclid and beyond.* Undergraduate Texts in Mathematics. Springer-Verlag, New York, 2000, pp. xii+526. ISBN: 0-387-98650-2. DOI: `10.1007/978-0-387-22676-7`.

[Her68] I. N. Herstein. *Noncommutative rings.* The Carus Mathematical Monographs, No. 15. Published by The Mathematical Association of America; distributed by John Wiley & Sons, Inc., New York, 1968, pp. xi+199.

[Mil80] James S. Milne. *Étale cohomology.* Princeton Mathematical Series 33. Princeton, N.J.: Princeton University Press, 1980. ISBN: 0-691-08238-3.

[Row91] Louis H. Rowen. *Ring theory.* Student. Academic Press, Inc., Boston, MA, 1991, pp. xxviii+623. ISBN: 0-12-599840-6.

[Ser77] Jean-Pierre Serre. *Linear representations of finite groups.* Translated from the second French edition by Leonard L. Scott, Graduate Texts in Mathematics, Vol. 42. Springer-Verlag, New York-Heidelberg, 1977, pp. x+170. ISBN: 0-387-90190-6.

Department of Math, Ben-Gurion University, Be'er-Sheva, Israel
Email address: `mailto:kamenskm@math.bgu.ac.il`
URL: `http://mkamensky.notlong.com`