A FLAVOUR OF NONCOMMUTATIVE ALGEBRA (PART 1)

VIPUL NAIK

Abstract. This is the first part of a short two-part write-up on noncommutative algebra. The material is related to course material covered by Professor Victor Ginzburg in Math 325 at the University of Chicago. The author would like to thank Emily Riehl and Michael P. Miller for their comments and suggestions.

1. Noncommutative rings: the basic definition

1.1. **Rings.** The definition of ring that we shall use in this write-up is given below. Note that this is not *the correct definition* in any absolute sense.

Definition (Ring)**.** A ring is a set R equipped with two binary operations $+$ (called addition) and $*$ (called multiplication) and constants 0 and 1 such that:

- R is an Abelian group under $+$, with additive identity 0 and inverse operation denoted by the prefix symbol $-$.
- The multiplicative operation of R admits 1 as a two-sided identity.
- The multiplication operation is associative.
- The multiplication operation satisfies left and right distributivity over addition (some people say it the other way around). In symbols:

$$\begin{aligned} a*(b+c) &= (a*b)+(a*c) \\ (a+b)*c &= (a*c)+(b*c) \end{aligned}$$

Lots of people have alternative views on what a "reasonable" degree of generality would be:

- People who work only in commutative algebra define a "ring" to be a "commutative ring" viz., a ring where the multiplication operation is commutative.
- In a number of applications it is important to relax the assumption that the ring should have a two-sided multiplicative identity; in fact, that happens even for commutative rings. Examples are the ring of functions under convolution, where there only exist "approximate identities".
- In some applications, the associativity assumptions are relaxed.
- There are applications where only *one* of the distributivity laws holds. A simple example is the set of all functions from an Abelian group to itself, where the addition is pointwise and the multiplication is by composition. The distributivity law holds only one way (of course, if one restricts to *endomorphisms* then both distributivity laws hold).

It is important to bear in mind that there's nothing canonical about the definition.

1.2. **Ring homomorphisms.** We assume that our ring homomorphisms preserve the binary operations $+$ and $*$ and send 0 and 1 to 0 and 1 respectively. The last part is important; in particular it means that the definition of subring requires that the subset should be an Abelian subgroup under $+$ and closed under $*$, and that it should contain the *same two-sided identity* which is there in the whole ring.

1.3. **Algebras over commutative rings.** Let R be a commutative ring. Then a ring A is said to have the structure of a R-algebra if we have a nonzero homomorphism from R to the *center* of A. It is particularly important to note that for an R-algebra we require that the image of R should be inside the center of the huge ring.[1]

[1] Again, conventions vary widely in what it means to be an "algebra" over a ring

A case of particular interest is when R is a field, in which case we denote it by $\mathbf{k}$. In this case, the map from $\mathbf{k}$ to the center of A is actually injective (because the kernel of the map must be trivial), and thus we can think of a copy of $\mathbf{k}$ sitting inside A.

Another case of particular interest is when $R = \mathbb{Z}$, because every ring A can be viewed in a unique way as a $\mathbb{Z}$-algebra: namely by the map sending 1 to 1.[2]

1.4. **A glimpse of the gameplan.** Noncommutative rings are multi-headed monsters – there are so many facets to them that in general it's hard where and how to get a handle on them. On the one hand, we have this huge monstrous ring and we want to study it globally, on the other hand we want to get friendly with the individual elements in the ring. These approaches are not wholly contradictory, because the addition and multiplication operations of the ring ensure that every element affects every other element.

Based on personal taste and comfort, I will begin first with an exploration of the "elements" of the ring. This is probably not the "correct" way: the correct way would emphasize looking at objects in their entirety. But playing around with elements can start off by giving us a feel of how manipulations can (and more importantly, cannot) be carried out in a noncommutative ring.

Again, based on personal taste, I will start out by defining some notions in the context of a set with a binary operation, and then study these notions specifically for a ring. This will help us appreciate what facets of ring structure get used.

After doing the "element-wise exploration", I shall study "ideals" and "modules". Again, there are two related methods of study: one is to view everything as modules, and the other is to focus on ideals as living set-theoretically inside the ring. Each approach gives its own share of insights, so I will compare and contrast them.

I will use two more "elementary" branches of algebra to provide parallels, analogy and motivation:

- Group theory
- Commutative algebra

2. Element-wise exploration

2.1. **A list of definitions and facts.** Here are some definitions involving a *single* binary operation $*$:

Definition.

- A **left identity element**(defined) for $*$ is an element e such that $e * a = a$ for all a; analogously define right identity element. An **identity element**(defined) for $*$ is an element which is both a left identity element and right identity.
- A **left nil element**(defined) for $*$ is an element n such that $n * a = n$ for all a; analogously define right nil element. A **nil element**(defined) is an element that is both a left nil element and a right nil element.
- An **idempotent**(defined) for $*$ is an element u such that $u * u = u$. Note that any left or right nil or any left or right identity element is an idempotent; there could be many other idempotents, though.
- A **left-cancellative element**(defined) for $*$ is an element a such that $a * x = a * y \implies x = y$. Analogously define **right-cancellative element**(defined). A **cancellative element**(defined) is an element that is both left-cancellative and right-cancellative.
- Suppose $*$ admits a two-sided identity e. A **left-invertible element**(defined) for $*$ is an element a for which there exists b such that $b * a = e$. Such a b is termed a left inverse for a. A **right-invertible element**(defined) for $*$ is an element a such that there exists a b such that $a * b = e$. Such a b is termed a right inverse for a. An **invertible element**(defined) is an element which has a left and right inverse *and* such that both are equal
- Suppose $*$ admits a nil element n. A **left nil divisor**(defined) for $*$ is an element a such that there exists a $b \neq n$ for which $a * b = n$. Analogously define right nil divisor.
- Suppose $*$ is associative as well. Then a **nilpotent**(defined) for $*$ is an element a such that $a^r = n$ for some positive integer r.

Now some easy-to-prove, but important-to-remember facts:

[2] In the language of category theory, $\mathbb{Z}$ is the initial object in the category of rings.

(1) If $*$ admits a left identity and a right identity, they are equal, and we hence get an identity element. The proof of this requires no additional assumptions on $*$.
(2) if $*$ admits a left nil and a right nil, they are equal, and we hence get a nil element.
(3) If $*$ admits an identity *and* is associative, any left-invertible element is left-cancellative and any right-invertible element is right-cancellative.
(4) If $*$ admits an identity *and* is associative, any left inverse and right inverse of an element must be equal. Hence an element that is both left-invertible and right-invertible must be invertible.
(5) If $*$ admits a nil and is associative, a left nil divisor cannot be left-cancellative, and a right nil divisor cannot be right-cancellative.
(6) If $*$ admits a nil and is associative, any nilpotent is both a left nil divisor and a right nil divisor.

So much for general binary operations. In the case of rings, the multiplication operation $*$ is associative, and admits a nil (namely 0)[3] as well as an identity (namely 1). For rings, we use the term "left zero divisor" for left nil divisor and "right zero divisor" for right nil divisor; a lot of the rest of the terminology remains the same.

For rings, we have some further nice things. For a general binary operation, we had seen that left nil divisors cannot be left-cancellative. For *nonzero* rings, the converse is true: any element is left-cancellative iff it isn't a left nil divisor. The idea behind the proof is to use distributivity to convert an equality of the form:

$$a * b = a * c$$

to an expression of the form:

$$a * (b - c) = 0$$

While this may seem very trivial, there is something deep going on here: namely the underlying *additive* structure is forcing a kind of uniformity on the multiplicative structure. We shall see this uniformity pop up at various places; for instance, when we try to define ideals, and look at the behaviour of individual elements. Suffice it to say that for a general binary operation $*$ there is no way to "move" all terms to one side; however, the fact that we have an underlying Abelian *group* structure allows us to rewrite any equation in the form something $= 0$.

2.2. **More of the gameplan.** In this section we shall explore interesting properties for elements in rings. The focus is largely on how well-behaved the *multiplicative* structure of the ring is, the additive structure is in any case fairly well-behaved. As we just saw:

> The multiplication operation of a ring is an associative binary operation with a (two-sided) identity element (called 1) and a (two-sided) nil or zero element (called 0). Thus, the multiplication operation of a ring gives it the structure of a *monoid with zero*. However, there are a lot of properties of the multiplication operation of a ring that are not shared by arbitrary monoids with zero.

While studying how nice the multiplicative structure of a particular ring is, we shall see that there are two particularly *nice* situations that could occur:

- Two elements *commute*. In other words, $a * b = b * a$. If any two elements commute, then the ring is commutative. Life is a lot simpler for commutative rings because we can group terms together the way we want.
- An element is *invertible*. Even though invertible elements may not commute with all elements, we can still *push them past* other elements by multiplying by their inverses. Other manipulations also become feasible.

We shall also see that there are two general trends that elements may have:

- The trend of being *close to* 1: Invertible elements are the closest to 1.
- The property of being *close to* 0: There are many elements that are good candidates for being close to 0. Possibly the closest to 0 are the nilpotents, after which come the zero divisors.

[3]The proof that 0 is a nil follows from distributivity

2.3. **Equations satisfied and subring.** Recall the following trivial observation we made a little while ago:

> In a ring, being left-cancellative is equivalent to *not* being a left-zero divisor, and being right-cancellative is equivalent to *not* being a right-zero divisor. Further, left-invertible implies left-cancellative and right-invertible implies right-cancellative.

Pictorially, we can think of the elements of the ring as divided into three classes:

(1) The left-invertible elements
(2) The left-zero divisors
(3) The elements that are neither left-invertible nor left-zero divisors, viz., the left-cancellative elements that are not left-invertible.

For commutative rings, the "left" can be dropped; for noncommutative rings we get analogoues by replacing "left" with "right".

In a finite ring, the third possibility cannot hold: every element is either left-invertible or a left-zero divisor. For infinite rings, we may in general have all the cases:

- In the polynomial ring $\mathbf{k}[x]$ there are no left-zero divisors other than 0. However, only the scalars are left-invertible. Thus all nonzero elements are of type (1) or (3)
- In the ring $\mathbf{k}[x]/(x^2)$ every element is either invertible or a zero divisor, viz., all elements are of type (1) or (2).
- In the ring $\mathbf{k}[x,y]/(x^2)$ the scalars are of type (1), the element x is of type (2) and the element y is of type (3).

We now make an elementary but deep observation. Suppose R is a subring of S. Then any left zero-divisor in R continues to be a left zero divisor in S; however, there may be elements which are not left zero divisors in R which become left zero divisors in S.

Similarly any left-invertible element in R continues to remain left-invertible in S, but there may be elements in R which become left-invertible in S.

The proof in both cases is the same, so we sketch it to observe what the proof does.

Claim. Let R be a subring of S. If a is a left-zero divisor in R, a is a left-zero divisor in S.

Proof. Since a is a left-zero divisor in R, there exists a $b \neq 0$ in R such that $a * b = 0$. Now since R is a subring of S, this b also lives in S, so we get $0 \neq b \in S$ such that $a * b = 0$. Hence a is a left-zero divisor in S. □

The crux of the proof is that being a left-zero divisor is dependent on the *existence of a witness inside the ring*, and enlarging the ring only increases the possibility of the existence of such a witness.

This might lead us to the question:

- Suppose a is an element of a ring R. Is there a bigger ring S containing R such that a becomes a left-zero divisor in S?
- Suppose a is an element of a ring R. Is there a bigger ring S containing R such that a becomes left-invertible in S?

The answer to this question is interesting. The *non-existence* of a witness for a being a left-zero divisor in R does not stop us from having a witness in S. However, the existence of a witness for a being *left-invertible* prevents a from being a left-zero divisor in any bigger ring. In other words, if a is already a left-invertible, there is no hope of its being a left-zero divisor in any bigger ring. Similarly, if a is already a left-zero divisor, there is no hope of its being left-invertible in any bigger ring.

In our language of classes (1), (2) and (3) we can put it this way:

> Let R be a subring of S. Then elements of type (1) (left-invertibles) in R continue to remain of type (1) in S. Elements of type (2) (left-zero divisors) in R continue to remain of type (2) in S. Elements of type (3) may remain of type (3) or may become of type (1) or (2) in S.

2.4. **The total quotient ring.** We concluded the last section by saying that an element that is neither left-invertible nor a left-zero divisor in R could in principle become either in S. But is it always possible? Here are some partial answers:

- When R is a commutative ring, there is a construction of a commutative ring $K(R)$, called the *total quotient ring* of R, such that $K(R)$ contains R and every element that is *not* a zero divisor for R becomes invertible in $K(R)$. The idea is to formally adjoin inverses for all elements that are not zero divisors, and then define addition and multiplication of fractions in the "usual" way.[4]
- When R is a commutative ring, it is also possible to construct a commutative ring S containing R such that every non-invertible element inside R becomes a zero divisor inside S (I am not completely sure of this construction, so verify it yourself).
- When R is a noncommutative ring, it may happen that $a \in R$ is not a left-zero divisor, but still cannot be made left-invertible in any bigger ring. In other words, there could be other "obstructions" to making the element left-invertible. An example is: a is *right*-invertible but not left-invertible. In this case, a cannot have a left-inverse in any bigger ring because any left-inverse and right-inverse are forced to be equal.

  I am not sure whether one can explicitly list *all the obstructions* to making an element left-invertible in a bigger ring.

2.5. **Nilpotents.** We saw that being invertible (left, right, or two-sidedly) and being a zero divisor were both *existential* conditions: they require the existence of another element in the whole ring satisfying certain equations. Thus, it is possible to enlarge the ring and make an element invertible even if it was not earlier so. Hence it is not a property truly intrinsic to the element.

On the other hand, the property we now talk of is a truly intrinsic property, namely that of being nilpotent. An element in a ring is termed nilpotent if some positive power of it is 0. Note that if R is a subring of S and a is an element of R, a is nilpotent inside R if and only if a is nilpotent inside S.

These facts are, however, not as comforting as they might appear. For commutative rings, being nilpotent is a powerful property, primarily because sums and products of nilpotents are still nilpotent.[5] For noncommutative rings, nilpotence is a very weak property, and one typically needs to replace it by a stronger condition called *strongly nilpotent*. Unfortunately this stronger condition now has a dependence on the ring at large and is not local enough. We shall discuss this later.

2.6. **Idempotents.** An **idempotent**(defined) in a ring R is an element e such that $e^2 = e$. A *nontrivial* idempotent is an idempotent other than 0 or 1. Nontrivial idempotents are superb examples of elements which refuse to be close to either 0 or 1; they strike out on their own.

Idempotents are closely related to splitting the ring as a direct sum; we shall discuss this later. For now, a little claim:

Claim. If e is an idempotent in a ring, so is $(1-e)$. Moreover, they commute, and their product is 0.

Proof. $e = e^2$ rearranges as $e(1-e) = 0$ and also as $(1-e)e = 0$, so they commute and their product is 0. Also, $(1-e)^2 = 1-e$ by a direct check. □

Idempotents are again not as comforting as they might appear: in general, a product of idempotents need not be idempotent. It *is* true that a product of *commuting* idempotents is still idempotent. We shall return to this theme later.

2.7. **Inner automorphisms and conjugate elements.**

Definition (Inner automorphism of a ring)**.** Let R be a ring (with identity element). If a is an element of R with two-sided inverse denoted a^{-1}, the **inner automorphism**(defined) induced by a is the map:

$$x \mapsto axa^{-1}$$

This is also called *conjugation* by a.

Claim. The inner automorphism induced by a, as in the above definition, is indeed an automorphism.

[4] One has to be a bit careful with this; it's not the "usual" way for those not used to it

[5] The proof for sums uses the "binomial theorem" which is applicable in any commutative ring.

Proof. We do not write out all the proof details, but indicate which axioms of the ring structure need to be used for each check:

- The fact that the map preserves the additive structure follows from a combination of left and right distributivity
- The fact that the map preserves the multiplicative structure follows from the associativity of multiplication, and the fact that a^{-1} is a *left* inverse of a
- The fact that the map sends 1 to 1 follows from the fact that a^{-1} is a *right* inverse of a
- The fact that the map is invertible follows from the fact that conjugation by a^{-1} inverts the map. This in turn relies on a^{-1} being a *two-sided inverse* for a

□

Some further checking shows that in fact if G is the group of elements with two-sided inverses in R (under multiplication) then the action by conjugation is a *group action* of G on R. We can thus define the notion of two elements of a ring being conjugate:

Definition (Conjugate elements)**.** Two elements x and y of a ring R are said to be **conjugate**(defined) if there exists an invertible element $a \in R$ such that $axa^{-1} = y$.

The fact that we have a group action tells us that the relation of being conjugate is an equivalence relation (because its the relation of being in the same orbit under a group action).

Note that if x and y are conjugate, then all ring-theoretic properties satisfied by x are also satisfied by y, because conjugation is in particular an inner automorphism. In fact because inner automorphisms of rings can be extended to bigger rings, it is true that any property that x may have in any bigger ring containing R, is also shared by y.[6]

So any conjugate to a left zero divisor is a left zero divisor, any conjugate to a left-invertible is left-invertible, any conjugate to a nilpotent is nilpotent, any conjugate to an idempotent is idempotent.

Here's an important definition:

Definition (Center of a ring)**.** The center of a ring R is the set of those elements $z \in R$ for which $zx = xz$ for all $x \in R$.

The center of a ring is a subring. It's *not* an ideal.

And now a few facts:

- Every element in the center of the ring is fixed by all inner automorphisms. The converse is not true in general. It is true that if the ring is *additively generated* by the invertible elements, then the center is precisely the set of elements fixed by all inner automorphisms. Rings generated additively by their invertible elements are of great interest, and we shall return to this theme later.
- If x and y are conjugate in R, and z is in the center of R, then $z+x$ and $z+y$ are also conjugate in R; in fact the same conjugating element will do. While this observation may seem innocuous, it is crucial while understanding spectra. We shall return to this later.

2.8. **How close are ab and ba?** If a is invertible, there is an inner automorphism (namely conjugation by a) sending ba to ab, so ab and ba behave in practically the same way. If b is invertible, there is an inner automorphism (namely conjugation by b) which sends ab to ba, and hence again their behaviour is the same.

But what happens if neither a nor b is invertible? Can one find some automorphism sending ab to ba? If not, can one still say that ab and ba behave in a reasonably similar manner?

The answer to the first question is *no*. In fact, it could happen that $ab = 0$ but $ba \neq 0$. The category theorist's example would be the quotient $\mathbf{k}\langle x, y\rangle / \langle\langle xy \rangle\rangle$ viz., the quotient of the free associative algebra in two variables by a one-way product being 0. A more "concrete" example is the matrix ring $M_2(\mathbf{k})$ with matrices:

[6] This is the fact, yet again, that inner automorphisms are extensible

$$a = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix}, b = \begin{pmatrix} 0 & 0 \\ 1 & 0 \end{pmatrix}$$

This may seem like a big blow, but there are bigger blows. Suppose we make the following definitions:

Definition (Weakly conjugate)**.** Two elements $x, y \in R$ are termed **weakly conjugate**(defined) if there exist $a, b \in R$ such that $ab = x$ and $ba = y$.

Then the relation of being weakly conjugate is reflexive and symmetric but it is *not* transitive.

On the plus side, though, we have the following:

- If x and y are weakly conjugate (viz., $x = ab$ and $y = ba$) then one is nilpotent if and only if the other is. Moreover the difference between their "nilpotence" (the smallest exponent which gives 0) is at most 1.

  The proof of this is elementary but yet again deep. Write $x = ab$ and $y = ba$. Then write x^r as:

$$abababab \dots ab$$

  r times. Clearly, if this is 0, then ba repeated $r + 1$ times is also 0.
- If x and y are weakly conjugate, and z is an invertible element in the center, then $x - z$ is invertible if and only if $y - z$ is invertible (this is a paraphrasing of the "famous" fact that $1 - ab$ is invertible if and only if $1 - ba$ is invertible).

The implication of this (for those who care) is that if R is a $\mathbf{k}$-algebra, and x and y are weakly conjugate in R, then the spectrum of x is the same as the spectrum of y.

3. Ideals

3.1. **Moving beyond elements.** Studying a huge and monstrous ring by looking at one element at a time has its limitations, so it's time to move to studying certain important subsets of rings that give a good idea of the ring structure. Unfortunately, there are a large number of approaches to this, and each approach has its own passionate "exponents". Since my own knowledge and intuition is better in group theory than in nocommutative ring theory, I begin by quickly reviewing what happens in group theory, and then closely studying the analogy with the noncommutative ring case.

3.2. **The group theory situation.** If we start off with a group, two interesting things we can do with it are:

- Study sets on which the group acts
- Study quotient groups of that group

Let's quickly recall what we know about "sets" on which a group acts (on the left). If a group acts on a set, the set can be partitioned into *orbits*. The group can then be viewed as acting on each orbit, and the action within each orbit is a *transitive* group action. Then the following are true (believe me, or prove it yourself, or check it out somewhere):

- For any point in the orbit, we can talk of its *isotropy subgroup* or *stabilizer*: the subgroup of the group which fixes that point. Then the isotropy subgroups at any two points in an orbit are conjugate subgroups.

  Let's see this with notation. Suppose the group is G, the orbit is S, and the two points are x and y. Suppose G_x and G_y are the isotropy subgroups respectively, and $g \in G$ is such that $g.x = y$, then:

$$gG_xg^{-1} = G_y$$

- In fact the set of isotropy subgroups is a conjugacy class of subgroups in the whole group.
- The group action on the orbit is completely determined by knowing the isotropy subgroup at any *one* point. Further, the action is the same as the action on the *coset space* of that subgroup.

A summary is the following:

> Transitive actions of groups on sets (upto equivalence) are in one-to-one correspondence with conjugacy classes of subgroups, where the map one way sends a transitive action to its set of isotropy subgroups and the map the other way sends a conjugacy class of subgroups to the action of the group on the left coset space of any one of those subgroups.

If we instead make groups act on the right, then we have the same result as above, except that the "left coset space" gets replaced by the "right coset space".

Another fact which is of interest is the following:

> For a transitive group action, the intersection of isotropy subgroups at all points in the set is a *normal* subgroup; it is the intersection of conjugates of any one isotropy subgroup, and can also be described as the kernel of the group homomorphism to the associated symmetric group on that set.
>
> Given a subgroup, the intersection of all its conjugates is sometimes termed its **normal core**(defined).

On a somewhat different note, we may be interested in studying surjective group homomorphisms from the group to other groups. In this case, we have the result:

> Surjective group homomorphisms emanating from the group (upto equivalence) are in one-to-one correspondence with normal subgroups, where the map one way sends a surjective group homomorphism to its kernel, and the map the other way sends a normal subgroup to the canonical quotient map by that normal subgroup.

3.3. **Modules: simple and cyclic.** Groups have very little structure, so one can make do with acting them on sets; rings, on the other hand, have more structure, so they deserve to act on something better than sets. We typically make rings act on Abelian groups, and an Abelian group endowed with a ring "acting" on it is termed a module. Thus group actions on sets are analogous to ring actions on Abelian groups.

Definitions I am not giving but expecting knowledge of: Left/right module, submodule, quotient module

The analogy stops being perfect from now onwards. Group actions are particularly nice; the whole set can be broken down into orbits. This is essentially because of the deep fact that every element in the group has an inverse. When rings act, however, we don't get "orbits" because a whole lot of elements of the ring don't have multiplicative inverses. If you're getting lost at this point (and even otherwise), it's time to backtrack a bit and see what's happening.

Definition (Accessibility)**.** Suppose R is a monoid with identity 1 (set with associative unital binary operation) and R acts on a set S, viz., there is a map $. : R \times S \to S$ such that $(r_1 r_2).s = r_1.(r_2.s)$ and $1.s = s$. Then for $s, t \in S$, say that t is **accessible**(defined) from s if there exists $r \in R$, such that $r.s = t$.

Claim. The relation of being accessible is reflexive and transitive but not necessarily symmetric. (It's a *quasi-order*, but that need not bother us right now).

Proof. Reflexivity follows from $1.s = s$ and transitivity follows from $r_1.(r_2.s) = (r_1 r_2).s$. Symmetry does not come for free because elements of the monoid need not be invertible. □

A ring action is in particular a monoid action of its multiplicative monoid, and ring actions have *no hope* of being transitive because a ring has an element called 0 which maps everything to 0. Thus, there is no hope of partitioning the module into orbits and then studying the ring action on each orbit.

One may argue that 0 is a pathology, so the closest one can get to a transitive action is if all the nonzero elements are accessible from each other. This "motivates" a definition:

Definition (Simple module)**.** Let R be a ring. A R-module M is termed **simple**(defined) if $M \neq 0$ and for any two elements $0 \neq m_1, m_2 \in M$, there exists $r \in R$ if $rm_1 = m_2$.

A somewhat weaker definition would state that there is *some* element from which every element is accessible.

Definition (Cyclic module)**.** Let R be a ring. A R-module M is termed **cyclic**(defined) if $M \neq 0$ and there exists $m \in M$ such that for every $m' \in M$ there is a r satisfying $r.m = m'$. Such a m is termed a cyclic element for M.

In a simple module every nonzero element is cyclic.

Now in group theory, once we had a transitive group action on a set, we used that to get hold of a subgroup by looking at the isotropy at a point. Ring theory, being more destructive, instead looks at the annihilator of an element. To cut a long story short, we first make some definitions:

Definition (Left ideal, right ideal, two-sided ideal)**.** Let R be a ring. A subset I of R is termed:

- A **left ideal**(defined) of R if I is a subgroup of R under $+$ and if $rx \in I$ for any $r \in R$ and $x \in I$ (note that the element from the *whole* ring is on the left).
- A **right ideal**(defined) of R if I is a subgroup of R under $+$ and if $xr \in I$ for any $x \in I$ and $r \in R$.
- A **two-sided ideal**(defined) of R if it is both a left and a right ideal.

Now a long list of facts, each easy to verify in its own right, but hard to come up with on one's own:

- R can be treated as a left module over itself, or alternatively as a right module over itself.
- With this view, the left ideals of R are precisely the left R-submodules of R, and the right ideals of R are precisely the right R-submodules of R.
- If I is a left R-module, then by the quotient module structure, R/I also becomes a left R-module.
- Suppose M is any left R-module. Define the **annihilator**(defined) of $m \in M$ as the set of $r \in R$ such that $r.m = 0$. Then the annihilator of m is a left ideal. Further the intersection of annihilators of *all* elements of M is a *two-sided ideal*.
- If M is a right R-module, the right annihilator can be defined analogously, and is now a right ideal. Again, the intersection of the annihilators of all elements is a two-sided ideal.
- If M is a *cyclic* left R-module, with cyclic element $m \in M$, then the action of R on M is equivalent to the action of R on R/I where I is the annihilator of m.
- Conversely, given a left ideal I, the quotient module R/I is a cyclic R-module. In fact, we get a bijection between left ideals and cyclic R-modules.

  Further the annihilator of this cyclic R-module can be described purely ring-theoretically as the largest two-sided ideal contained inside I. The largest two-sided ideal contained inside a left ideal is termed its **right core**(defined).
- A module over a ring is termed **faithful**(defined) if the annihilator of the module is 0.

So if one were to try setting up a correspondence between group actions on sets and ring actions on modules, it might look something like this:

Group notion	Ring notion
Group	Ring
Set with left group action	Left module over ring
Subgroup	Left ideal
Isotropy subgroup	Annihilator
Transitive group action	Cyclic module
Kernel of group action	Annihilator of module
Normal subgroup	Two-sided ideal
Normal core of a subgroup	Right-core of a left ideal
Faithful group action	Faithful module

The analogies will keep growing both in depth and in size as we proceed, as will the stark contrasts.

3.4. **Intersection of conjugates?** When a group acts transitively on a set, all points of the set look similar, and hence the fact that the isotropy subgroups at all points are conjugate is not surprising.

However, for a cyclic module over a ring, all the points in general do *not* look similar, so the annihilators at different points may not look very similar. We can still characterize the annihilator of any element in terms of the annihilator of the cyclic element as follows:

Definition (Division notation)**.** Let I be a left ideal in a ring R and $x \in R$. The notation $[I : x]$ is understood to refer to the set of those elements $r \in R$ such that $rx \in I$. Note that $[I : x]$ is also a left ideal.

We now observe the following:

Definition. Suppose M is a cyclic R-module with cyclic element m. Let I be the annihilator of m. Then the annihilator of $x.m$ where $x \in R$, is $[I : x]$.

Thus the annihilator of the whole module M, or equivalently the "right-core" of the ideal I is given by:

$$\bigcap_{x \in R} [I : x]$$

Contrast this with the expression for the kernel of a group action with isotropy subgroup H:

$$\bigcap_{g \in G} gHg^{-1}$$

3.5. **Commutativity and invertibility: both are good.** Let's take a moment to relax and see what happens for a *commutative* ring. When the ring is commutative, any left ideal is in fact a two-sided ideal, and if M is a cyclic R-module, with I the annihilator of a cyclic element, then I annihilates M. In the notation above, each $[I : x]$ contains I, and hence the intersection of all of these *is* I.

Thus commutativity is good; it simplifies things. Invertibility is also good, as we see in the following claim.

Claim. Suppose R is additively generated by its invertible elements. Then for any left ideal I, its right-core is the intersection of all "conjugates" of I where by a conjugate of I we mean a left ideal obtained as rIr^{-1} for r an invertible element of R.

The key ingredient of the proof is to observe that instead of intersecting $[I : x]$ over all x, it suffices to intersect $[I : x]$ over a set of x which *additively generate* R. Slickly put, this is the identity:

$$[I : x] \cap [I : y] \subset [I : (x + y)]$$

Thus invertibility is also a "good" condition.

We shall come back to these themes again, in due course.

3.6. **Primitive rings.** This is the first place where ring theory departs seriously from group theory. There *is* a parallel between so-called primitive group actions and primitive ring actions; however, this parallel is not what one would naively expect from the mechanism by which we have set up the analogy so far. Since our concern for the moment is rings, we will not go into what primitive group actions are, and instead focus on the notion of primitive rings.

We had set up a correspondence between left ideals and *cyclic* modules over the ring. Under this correspondence it turns out that *maximal* left ideals correspond to *simple* modules. To "see" the proof, perhaps we should redefine simple modules:

Alternative Definition (Simple module)**.** Let R be a ring, and M a R-module. We say that M is **simple**(defined) if M has no proper nonzero R-submodule.

I might put in an intuitive and/or formal explanation of the correspondence here later; for now, I move on to further stuff.

Definition (Primitive ring)**.** A ring is said to be **left-primitive**(defined) if it has a faithful simple left module. If by convention we are working with left modules, then we use the word "primitive" instead of left-primitive. Note that left-primitive is *not* the same as right-primitive.

Alternatively a ring is left-primitive if it has a maximal left ideal whose right core is trivial.

Primitive rings are in some sense analogous to "primitive groups" (a group is primitive if it has a maximal subgroup whose normal core is trivial), but this is of no interest to us right now.

3.7. **Two-sided ideals and the quotient ring.** We don't give the proofs this time, but state the result:

> Given any surjective ring homomorphism emanating from a ring, the inverse image of zero under it is a two-sided ideal. Further, this establishes a correspondence between two-sided ideals inside the ring and surjective ring homomorphisms emanating from it.

The significance of this in the context of ring actions is as follows. When a ring R acts on an Abelian group M, the action factors through an action of R_M where R_M is the quotient of R by the two-sided ideal which annihilates M. Moreover R_M acts faithfully.

This is analogous to the situation of a group action: quotienting out by the kernel of the group action gives a new group which acts faithfully.

3.8. **Simple ring.**

Definition (Simple ring)**.** A ring is said to be **simple**(defined) if it has no proper nonzero two-sided ideals.

Simple rings are the analogues of "simple groups".

Claim. Every simple ring is primitive.

Proof. The proof of this claim relies on a well-known fact: every left ideal of a ring is contained in a maximal left ideal (analogous statements hold for right ideals and two-sided ideals). Thus, given a simple ring, we can find a maximal left ideal containing the 0 ideal. The right core of this ideal is forced to be 0, since there are no other proper two-sided ideals. We have thus produced a maximal left ideal whose right core is 0, and thus the ring is primitive. □

In the commutative case the definitions of simple and primitive coincide, and the only simple rings are fields.

3.9. **Wrapping up this section.** In this section, a lot of diverse material was discussed, most of which we didn't understand too well. This is not surprising: most of this isn't understood too well by anybody. I used group theory to motivate, but now is the time to forget the group-theoretic parallels and review what we have done, purely ring-theoretically.

There are in general, two flavours:

- One flavour concentrates on everything as subsets of this ring. This is the language of left ideals, right ideals, two-sided ideals, right cores, intersections, conjugates etc.
- The other flavour tries to view everything in terms of a ring acting on something *outside*, viz., a module.

The advantage of the first flavour is that one sees everything happening within the *set* of the ring. The advantage of the second approach is that it may be a lot cleaner because we're not cluttering the interior of the ring with everything, we're looking at things from a more far-out perspective. We must be versatile enough to translate freely between the two approaches.

I first review everything done so far from the viewpoint of the first approach:

(1) Left (resp. right) ideals are subsets of a ring that are Abelian groups under addition and are closed with respect to left (resp. right) multiplication by arbitrary ring elements
(2) A two-sided ideal is a subset which is an Abelian group and is closed under left as well as right multiplication by arbitrary elements of the ring
(3) Given a left ideal, its "right core" is the largest two-sided ideal contained inside it. Analogously we can define the "left core" of a right ideal.
(4) For a commutative ring, left ideal, right ideal and two-sided ideal are equivalent notions.
(5) When the ring is additively generated by its invertibles, the right core of a left ideal is simply the intersection of all its conjugates. In general, that picture might break down.
(6) A left-primitive ring is defined as a ring with a maximal left ideal whose right core is trivial.
(7) A simple ring is defined as a ring with no proper nonzero two-sided ideals. Every simple ring is primitive.

And now for the module-based picture:

(1) There is a correspondence between cyclic left modules of a ring with *chosen cyclic element* and left ideals. The correspondence sends the cyclic left module to the annihilator of the chosen cyclic element. For the other way around, we make the ring act on the quotient space by the ideal.

    Given a cyclic left module, there may actually be multiple left ideals it corresponds to: the annihilators of its various cyclic elements. *These left ideals are not necessarily conjugate*, but they satisfy a weaker equivalent of being conjugate (that's called *being associate*, but LNGITRN).
(2) Under this correspondence, simple left modules correspond to maximal left ideals.
(3) The annihilator of any module is the intersection of the annihilators of all its elements. For a *cyclic* module, it is the right core of the annihilator of any cyclic element. For a *simple* module, it is the right core of the annihilator of *any* element.
(4) A module is called faithful if its annihilator is zero. Given any module over a ring, the module can be viewed as a faithful module over the quotient ring by the annihilator. This is the *effective* ring (the effective ring for R acting on M is sometimes denoted R_M).
(5) A left-primitive ring is a ring with a faithful simple left module.

4. Some murkier exploration

4.1. **Characteristic and normal.** This subsection is motivational, one can plunge straight into the next subsection without loss of continuity.

I introduce a few definitions:

Definition (Characteristic subset, normal subset)**.** A subset of a ring is termed **normal**(defined) if it is invariant under all inner automorphisms. A subset of a ring is termed **characteristic**(defined) if it is invariant under all automorphisms.

For commutative rings, every subset is normal. Our focus is on noncommutative rings. We are now interested in subsets which are left ideals. In general we have the following:

> Any characteristic left ideal is normal, and any two-sided ideal is also normal. However, a characteristic left ideal need not in general be two-sided, and a two-sided ideal need not in general be characteristic.

An interesting question is: under what conditions can we guarantee that every normal left ideal is two-sided? *A priori* the condition of being two-sided is much stronger than being normal; I is normal in R if for every invertible $x \in R$ we have:

$$xIx^{-1} \subset I$$

On the other hand for I to be two-sided, we have the much stronger requirement that for any two elements $x, y \in R$:

$$xIy \subset I$$

Again, we find that if the ring is additively generated by its invertible elements, then the two conditions become equivalent. In other words, for rings which have *enough* invertibles, normal left ideal is the same as two-sided ideal, and in particular we see that any characteristic left ideal must be two-sided.

Now if you're like me, you would probably say: *the condition of being additively generated by invertibles seems to occur very often*. Further, it in some sense yields similar conclusions to the condition that the ring be commutative. This is back to the old theme: both commutativity and invertibity are powerful conditions. When we have two powerful conditions both of which yield similar conclusions we try to replace them both by a weaker condition which still yields similar conclusions. This is what we do in the next subsection.

4.2. **Invariant elements.** I had earlier called these elements "weakly central", but it turns out that Cohn[7] uses the term "invariant element" for these, so I'll use the same term. I am not sure how standard this term is:

Definition (Invariant element)**.** An element x in a ring R is termed **invariant**(defined) if $Rx = xR$, viz., the left and right ideals it generates are the same.

One way of viewing this definition is that even if x doesn't commute with every element of R, every element of the form rx can be written as xr' for some r'.

Now some quick observations about invariant elements:

- Every central element (element in the center) is invariant. In particular, for a commutative ring, every element is invariant.
- Every invertible element is invariant. In fact if x is invertible we have $Rx = xR = R$.
- The set of invariant elements is multiplicatively closed, but it need not be closed under addition.

Observe that for a invariant element the left ideal generated by it is a two-sided ideal. This leads to a Schur's lemma-type statement:

Claim. In a simple ring, the only invariant elements are 0 and the invertible elements.

The classical example to keep in mind is the matrix ring over a field.

4.3. **The headache of left versus right.** What makes noncommutative algebra hard? There are many answers possible but perhaps the one that wins for me is that there's so much of a left-right headache. Commutativity eliminates this headache completely.

Even though groups are in general far from Abelian, we don't have distinct notions of "left subgroup" or "right subgroup": the reason is *invertibility*, which allows us to push stuff past stuff, possibly with a bit of change.

In general, for a noncommutative ring, there could be "left" properties it satisfies whose "right" analogues it fails to satisfy. On the other hand, there are some rare properties that are defined using a "left" language but turn out to be left-right symmetric; this is more the exception than the rule. So as not to lose sight of this headache, there is a useful notion of *opposite ring*:

Definition (Opposite ring)**.** The **opposite ring**(defined) to a ring R, denoted R^{op}, is a ring which as a set is identified with R, with the addition being the same, but with the multiplication defined as:

$$a * b = ba$$

[7] "Further Algebra and Applications" by P.M. Cohn; a veritable treasure-trove of all aspects of algebra

Actually the "opposite" can be defined more generally for any binary operation as reversing the order of the inputs. The left-right symmetry in the axiomatization of the ring ensures that the opposite of the multiplication still satisfies the required axioms.

In general, a binary operation may behave very differently from its "opposite" but there are some special cases where the behaviour is the same:

- Every group is isomorphic to its "opposite" group. The map is $x \mapsto x^{-1}$. However, when G is the group of invertibles of a ring R, the above map from G to G^{op} may not extend to a ring homomorphism from R to R^{op}, because it may not be additive.
- Every *commutative* ring is isomorphic (as a *ring*) to its opposite ring. The map in this case is the identity map.
- A matrix ring over a commutative ring is isomorphic (as a *ring*) to its opposite ring. The map in this case is the *transpose* map.

To say these with a little less verbiage, I'll introduce a term:

Definition (Self-opposite ring)**.** A ring is termed **self-opposite**(defined) if it is isomorphic (as a ring) to its opposite ring.

There's a stronger notion of an *involutory* ring where we actually require the isomorphism to be an involution. Matrix rings over commutative rings are involutory rings.

Self-opposite rings are *globally* nice: for instance, in a self-opposite ring, if there exist left zero divisors, there also exist right zero divisors; if there exist left ideals satisfying some property, there also exist right ideals satisfying the analogous "right" version of that property. But locally they could be badly behaved; one could still have left zero divisors which are not right zero divisors, or left-invertibles which are not right invertibles. For "local" good behaviour one needs conditions like being central, invertible, or invariant.

5. Nilpotents and maximals

5.1. **In the commutative case.** As usual, the complexity of noncommutative algebra is breathtaking, so it seems hard to give an intuition into it without using an analogy with commutative algebra. Thus I quickly review basic facts in commutative algebra before plunging into the noncommutative story.[8]

- In a commutative ring, a sum of nilpotents is nilpotent, and the product of any nilpotent with any element of the ring is a nilpotent. Both proofs rely heavily on being able to switch the order of terms in a product.
- The upshot of the above is that the set of nilpotent elements in a commutative ring is an ideal.[9] This ideal is termed the **nilradical**(defined) of the commutative ring.
- Given a maximal ideal, the quotient of the ring by the maximal ideal is a field. Thus, there can't be any nilpotents outside a maximal ideal, and hence every nilpotent element is contained inside every maximal ideal. In other words, the nilradical is contained inside every maximal ideal.
- To put this even more nicely, define the **Jacobson radical**(defined) of a commutative ring as the intersection of all its maximal ideals. Then the nilradical is contained inside the Jacobson radical. The two radicals need not in general be equal.
- If R is a subring of S, the nilradical of R equals the intersection of R with the nilradical of S. This traces back all the way to our earlier discussion of nilpotents: being nilpotent is a property that depends on the element alone and doesn't depend on the appearance of the huge ring in which it sits.

Things turn nasty in the noncommutative case; the nilpotent elements need not in general form an ideal. A sum of nilpotents need not be nilpotent, and a product of a nilpotent with an arbitrary element need not be nilpotent. One doesn't even need to be pathological to produce examples: the matrix ring over a field contains lots of nilpotents, and in fact the two-sided ideal generated by nilpotents is the whole ring.

There's a small consolation, though: the matrix ring is not *additively* generated by its nilpotents. One way of observing this is that every nilpotent matrix has trace 0, and since the trace is additive, the

[8]Those unfamiliar with the commutative story have the advantage of going over the same new stuff twice.

[9]We're commutative right now, so there is no need for saying "left" or "right" or "two-sided"

additive subgroup generated by nilpotents comprises only elements of trace zero. This consolation might appear small but in fact it's very significant.

5.2. **Nil ideals.** We now give a definition suitable for the noncommutative case:

Definition (Nilpotent ideal)**.** A left ideal I in a ring R is termed **nilpotent**(defined) if there exists some n such that any product of length n, of elements from I, is 0. In fancy language, $I^n = 0$.

It now turns out that a sum of two nilpotent left ideals is again nilpotent (an infinite sum need not be nilpotent). This also motivates a new definition:

Definition (Strongly nilpotent element)**.** An element a in a ring R is termed **strongly nilpotent**(defined) if any sequence of the form c_n, where $c_{n+1} \in c_n R c_n$ and $c_1 = a$, eventually becomes 0 (there need not be a fixed n that works for all sequences).

All elements in a nilpotent left/right ideal are strongly nilpotent.

Now the property of being nilpotent was a property truly intrinsic to an element: it didn't depend on the ring at large. The property of being strongly nilpotent, on the other hand, depends, via universal quantification on the ring at large. Thus, enlarging the ring might actually turn an element that was earlier strongly nilpotent, into an element that is *not* strongly nilpotent.

For instance, any strictly upper triangular matrix is strongly nilpotent in the ring of upper triangular matrices over a field, but it ceases to be strongly nilpotent in the ring of all matrices.

On the other hand, an element which is strongly nilpotent in a bigger ring will continue to be strongly nilpotent in a smaller ring.

Let's now tie up invariant with strongly nilpotent.

Claim. Any invariant nilpotent element is strongly nilpotent.

Proof. Suppose R is the ring and a is invariant, $a^n = 0$. Then any product $c_1 a c_2 a \dots c_n a$ can be written as $a^n c_1' c_2' \dots c_n'$ by repeated use of the property $aR = Ra$. Using $a^n = 0$, we see that the product is 0. With some thought, we see that this implies that a is strongly nilpotent. □

For noncommutative rings, there is no well-defined notion of nilradical. One definition is as the set of strongly nilpotent elements, but there are many other definitions.

5.3. **The Jacobson radical.** We give a quick definition of the Jacobson radical as many equivalent things, but we don't prove why they are equivalent:

Definition (Jacobson radical)**.** The Jacobson radical of a ring is defined in the following equivalent ways:

(1) It is the intersection of annihilators of all simple left modules
(2) It is the intersection of all maximal left ideals
(3) It is the set of a such that $1 + xa$ is invertible for every x in the ring
(4) It is the intersection of annihilators of all simple right modules
(5) It is the intersection of all maximal right ideals
(6) It is the set of a such that $1 + ax$ is invertible for every x in the ring
(7) It is the set of a such that $1 + xay$ is invertible for every x and y in the ring

The equivalence of (1), (2) and (3) is "easy", the equivalence of (4), (5) and (6) is equally "easy", but to relate them we need to go via (7).

Note that the Jacobson radical is *not* defined in the way one would naively suspect: as the intersection of all maximal two-sided ideals.

For commutative rings, we saw that the nilradical is contained inside the Jacobson radical, or equivalently that any nilpotent element is contained inside the Jacobson radical. Here, we have a similar result: any *strongly* nilpotent element is contained inside the Jacobson radical. Another way of putting this is that any nilpotent ideal is contained inside the Jacobson radical.

5.4. **Time to put restricting assumptions.** We've seen that things have already grown to staggering proportions of complexity. This wasn't surprising: we discarded the two assumptions that were most responsible for good behaviour: commutativity and invertibility. To get some measure of sanity, it is time to put a restrictive assumption on the kind of rings and modules that we are dealing with. The winning assumption is Artinianness.

Definition (Artinian module, Artinian ring)**.** A left module over a ring is termed **Artinian**(defined) if any descending chain of submodules stabilizes after a finite stage. A ring is termed **left-Artinian**(defined) if it is Artinian as a left module over itself, viz., any descending chain of left ideals stabilizes.

6. Local rings

6.1. **Commutative local rings.** Time to take a break again, and look at some particularly nice kind of rings: local rings. We first begin with the commutative case, where we see a number of examples.

Definition (Commutative local ring)**.** A commutative ring is said to be **local**(defined) if the set of non-invertible elements is an ideal.

Clearly this ideal is the unique maximal ideal of the commutative ring, and the quotient by that is a field, which is termed the **residue class field**(defined). Here are some examples:

- Consider the ring of formal power series in one variable t over a field $\mathbf{k}$, denoted $\mathbf{k}[[t]]$. This is a commutative local ring with the unique maximal ideal generated by t. The residue field is $\mathbf{k}$ itself.
- In the finite situation, consider $\mathbf{k}[t]/(t^n)$. This is a commutative local ring with unique maximal ideal generated by t. The residue field is again $\mathbf{k}$ itself.
- Consider the ring $\mathbb{Z}/p^n\mathbb{Z}$. This is a local ring with maximal ideal being the multiples of p. The residue field is the field of p elements.

The first two cases are qualitatively different from the third. In the first two cases, the whole ring can be viewed as an *algebra* over its residue field: there is an embedding of the residue field inside the whole ring. However, in the third case this can't be done: the ring $\mathbb{Z}/p^n\mathbb{Z}$ is not an algebra over the field of p elements.

All these cases are, however, very special in some respects, which need not generalize to arbitrary local rings:

- In all cases, the unique maximal ideal is principal: it is generated by a single element. In fact these are local principal ideal rings (almost like discrete valuation rings).
- In the second and third case, the unique maximal ideal is nilpotent. This is essentially because of a finiteness assumption. We shall soon see that for a local Artinian ring, the unique maximal ideal is nilpotent. In the first case, the unique maximal ideal is far from nilpotent; in fact it has no zero divisors either.

6.2. **Local rings in the noncommutative setting.** A good generalization of the definition of local ring to the noncommutative setting is given below:

Definition (Local ring)**.** A ring is said to be **local**(defined) if the set of non-invertible elements forms a two-sided ideal.

This definition forces that the set of non-invertible elements is:

- The unique maximal two-sided ideal
- The unique maximal left ideal
- The unique maximal right ideal
- The Jacobson radical

Further, the quotient of the ring by this ideal is a *division ring*. Time to define a division ring:

Definition (Division ring)**.** A **division ring**(defined) is a ring which has no proper nontrivial left ideals, and no proper nontrivial right ideals. In other words, every nonzero element is invertible.

A division ring can thus also be viewed as a special case of a local ring where the unique maximal ideal is 0.

Let's see the chain of implications between ring properties as seen so far:

- Field $\implies$ Division ring $\implies$ Local ring
- Division ring $\implies$ Simple ring $\implies$ Left-primitive ring

Being a division ring is stronger than being a simple ring as well as being a local ring. In fact, a ring which is both a simple ring and a local ring is a division ring (the unique maximal ideal is forced to be 0 by simplicity).

6.3. **Schur's lemma: an elementary form.** I'll assume knowledge of what "direct sum" of *two* modules means. I'll also assume knowledge of what endomorphism of a module means.

Before claiming Schur's lemma, let's try to examine what the endomorphism ring of a module looks like. The idea is that endomorphisms of a modules are *elements* of the endomorphism ring, and their properties as *endomorphisms* can be related to their properties as *ring elements*.

Let's first make a somewhat trivial claim.

Claim. An endomorphism of a left module over a ring is invertible as an element of the endomorphism ring, iff it is an automorphism.

The proof is direct.

An interesting question is: under what conditions does an endomorphism of a left module possess a left inverse, and under what conditions does it possess a right inverse? In general, modules may possess endomorphisms with one-sided inverses. An example is a countable direct sum of copies of the ring; we have an endomorphism that shifts each coordinate one place to the right, and an endomorphism that shifts each coordinate one place to the left. The composite of these is the identity only one way.

However, for reasonably "nice" modules, we expect that endomorphisms that have left inverses also have right inverses; indeed we shall see that for nice modules, much stronger conditions hold. Let's begin with the nicest modules.

Theorem 1 (Schur's lemma)**.** The endomorphism ring of a simple left module over any ring, is a division ring.

Proof. By the previous claim, it suffices to show that every nonzero endomorphism is an automorphism.

Suppose R is the ring and M is a R-module. Then if $a : M \to M$ is a nonzero R-endomorphism, the kernel of a is a R-submodule of M, and the image of a is again a R-submodule of M. Since a is nonzero, the kernel cannot be the whole of M, hence it is zero; since a is nonzero, the image cannot be zero, so it must be the whole of M. So the map is actually a bijection from M to M; hence it is an automorphism[10] □

[10]For modules over rings, an endomorphism which is a bijection is an automorphism. In category-theoretic language, the foregetful functor from R-modules to sets is "conservative"

6.4. **Length of a module.**

Definition (Module of finite length)**.** A left module M over a ring is said to have finite length if there exists an ascending chain of submodules $0 = M_0 \leq M_1 \leq M_2 \leq \ldots \leq M_n = M$ such that each M_i/M_{i-1} is simple.

It is true that given a module of finite length, any two such ascending chains have the same set of simple quotients occurring (upto isomorphism). This result is called the Jordan-Holder theorem. More specifically, the number of times each simple quotient occurs in one such ascending chain, equals the number of times it occurs in any other ascending chain.

Here are some interesting facts:

- A direct sum of two modules of finite length again has finite length. The idea is that we can combine the two ascending chains for each of the direct summands. In fact the length of the direct sum is the sum of the lengths of each of the summands.
- Any module of finite length is finitely generated. This follows from the fact that every simple module is cyclic, and we can use the generators for each of the cyclic quotients to get a generating set for the whole module.

  The converse is not necessarily true. In fact, the ring itself may not be of finite length as a module over itself: a typical example is $\mathbb{Z}$.
- Something even better is true for generating sets that we recover in the above fashion. Namely, suppose $0 = M_0 \leq M_1 \leq M_2 \leq \ldots M_n$ is an ascending chain with simple quotients, and we pick a generating set m_i with the propery that $m_i + M_{i-1} = M_i$. Then the action of any element $a \in R$ sends m_i to a R-linear combination of $m_1, m_2, \ldots, m_i$.

A good example to keep in mind is the case where R is the ring of upper trianglar matrices of order n over a field $\mathbf{k}$ and M is the vector space $\mathbf{k}^n$. Then M is an indecomposable module of finite length over R. The ascending chain of submodules in this case is as follows: M_i is the linear span of the first i basis vectors of $\mathbf{k}$. Each quotient M_i/M_{i-1} is the field $\mathbf{k}$ itself, which is certainly simple as a R-module. The elements m_i of the above description can be taken as the standard basis vectors.

6.5. **Indecomposable left modules.** An overdue definition:

Definition (Indecomposable left module)**.** A left module over a ring is termed **indecomposable**(defined) if it cannot be expressed as the direct sum of two nonzero submodules.

There is a curious connection between decompositions and idempotents. For this note that if M is a direct sum of submodules M_1 and M_2 then there are projection maps $p_1 : M \to M_1$ and $p_2 : M \to M_2$ which thus give *idempotent* endomorphisms of M. If M_1 and M_2 are nonzero, these are both nontrivial idempotents, their sum is 1 and their product is 0 (they are of the form e and $1 - e$). Note also that e and $1 - e$ commute.

Let's make a definition:

Definition (Indecomposable ring)**.** A ring is said to be **indecomposable**(defined) if if it has no non-trivial idempotents.

And now a claim.

Claim. The endomorphism ring of a left module is an indecomposable ring if and only if the left module is an indecomposable left module.

The idea behind the proof is to show that the idempotents are "projection operators" and yield a direct sum decomposition.

And another claim:

Claim. A ring is indecomposable iff it is indecomposable as a left module over itself, iff it is indecomposable as a right module over itself.

The idea is to look at what the projection operator does to 1. This is a curious example of a situation where the left and right notions coincide.

6.6. **Local rings are indecomposable.** A local ring has no nontrivial idempotents: to see this, note that if e is an idempotent, either e or $1-e$ lives outside the maximal ideal, hence one of them is invertible. But the only invertible idempotent is 1, and we are done.

One interesting question is: under what conditions can we guarantee that the endomorphism ring of a given indecomposable module is local? Let's answer this question by definition first:

Definition (Strongly indecomposable left module)**.** A left module over a ring is termed **strongly indecomposable**(defined) if its endomorphism ring is a local ring.

Strongly indecomposable left modules are indecomposable because local rings are indecomposable.

A lot of the things we see in coming sections can be generalized to strongly indecomposable left modules, but we shall refrain from the generalization here.

6.7. **Indecomposables of finite length.** Question: *Why are we looking at endomorphism rings at all?* The idea is that if the endomorphism ring of an indecomposable is sufficiently nice, then that module will be sufficiently well-behaved. We already know the nicest fact: the endomorphism ring of a simple module is a division ring (this is Schur's lemma, which has always been around in the air).

Weakening the condition of being simple to being indecomposable certainly leads to quite a loss of "simplicity" in the corresponding endomorphism ring, but we still know that for any indecomposable module, the endomorphism ring is indecomposable, which is a start. However, we would like conditions where it is a local ring, or something even better. Indeed it turns out that if we restrict to indecomposables of finite length, then their endomorphism rings are local rings, and more:

Definition (Completely primary ring)**.** A ring is said to be **completely primary**(defined) if every element in it is either invertible or nilpotent.

Division rings are completely primary: the only subtle difference being that 0 is the only nilpotent around.

Completely primary rings are local: the nilpotent elements form a two-sided ideal (this requires a bit of argument and playing around, but it's true). We're now in a position to state the big result.

Theorem 2 (Fitting's lemma)**.** The endomorphism ring of an indecomposable module of finite length is a completely primary ring.

Fitting's lemma is thus an analogue of Schur's lemma where "division ring" is weakened to "completely primary ring" and "simple module" is weakened to "indecomposable module of finite length".

To summarize, we make a little table:

Module property	Implied endomorphism ring property	Proof
Simple	Division ring	Schur's lemma
Indecomposable of finite length	Completely primary ring	Fitting's lemma
Strongly indecomposable	Local ring	(By definition)
Indecomposable	Indecomposable ring	(proof sketched above)

6.8. **Isotypical modules of finite length.**

Definition (Isotypical module)**.** A left module over a ring is termed **isotypical**(defined) if it is a direct sum of pairwise isomorphic simple modules. By isotypical of finite length, we mean that it is isotypical and has finite length; which is equivalent to saying that it is a direct sum of finitely many copies of a simple module.

For an isotypical module, there may not in general be a *unique* choice of simple modules into which to decompose it; for instance, the vector space $\mathbf{k}^n$ is an isotypical module over $\mathbf{k}$, but different choices of bases may yield different decompositions of it as a sum of simple modules (read: one-dimensional vector spaces) over $\mathbf{k}$.

However, once we have chosen a direct sum decomposition, we can study how endomorphisms of the module look. Suppose $M = S_1 \oplus S_2 \oplus \ldots \oplus S_n$ where each S_i is a simple module over the base ring R, and all the S_is are isomorphic.

Then to specify a R-linear map from M to M yields a R-linear map from each S_i to each S_j, as follows:

$$S_i \to M \to M \to S_j$$

where the left map is the inclusion and the right map is the projection. Conversely given a R-linear map from each S_i to each S_j, there's a unique way of gluing them together to get a R-linear map from M to M.[11]

Now the homomorphisms from S_i to S_j are parametrized by elements of a division ring (because S_i is isomorphic to S_j) and we can thus write endomorphisms of M as matrices of order n whose entries come from the division ring which happens to be the endomorphism ring of the simple module S_1 (this matrix representation depends on first choosing an isomorphism between S_1 and each of the S_js). Miraculously once we write down the matrix, it turns out that multiplying the matrices is equivalent to composing the endomorphisms, and we see that:

> The endomorphism ring of an isotypical module of finite length is a matrix ring over a division ring, where the division ring is the endomorphism ring of the underlying simple module, and the order of the matrix ring is the number of copies of the simple module in the isotypical module.

6.9. **Semisimple modules of finite length.** We state another facet of Schur's lemma:

Theorem 3 (Schur's lemma part 2)**.** Let M and N be non-isomorphic simple R-modules. Then there is no nonzero R-homomorphism from M to N.

The idea is the same as that of the previous proof: any nonzero R-homomorphism must have trivial kernel and full image, but that would make M and N isomorphic.

We now define semisimple modules of finite length:

Definition (Semisimple module of finite length)**.** A **semisimple module of finite length**(defined) over a ring is a left module which can be expressed as a direct sum of finitely many simple modules.

[11] For the category theorist: here we are using that finite direct sums are both a product and a coproduct in the category of modules over rings. The inclusion map encodes the fact that it's a coproduct and the projection map encodes the fact that it's a product

Now what's happening is as follows. The particular decomposition of a semisimple module as a direct sum of simple modules need not be unique; we saw that even for isotypical modules it isn't unique. However, the following are true:

- Given a semisimple module and a direct sum decomposition into simple modules, club together all the simple modules of a particular isomorphism type. Then the direct sum of those is termed the **isotypical component**(defined) of the module for that isomorphism type. The isotypical component for a particular isomorphism type is independent of the choice of decomposition.
- Any endomorphism preserves the isotypical components, viz., it sends each isotypical component to within itself.
- The endomorphism ring is a direct sum of endomorphism rings of each of the isotypical components. Each of these is a matrix ring over a division ring, so *the endomorphism ring of a semisimple module of finite length is a direct sum of matrix rings over division rings.*

We are now in a position to considerably enlarge our table relating properties of modules to properties of their endomorphism rings:

Module property	Implied endomorphism ring property	Proof
Simple	Division ring	Schur's lemma
Indecomposable of finite length	Completely primary ring	Fitting's lemma
Strongly indecomposable	Local ring	(By definition)
Indecomposable	Indecomposable ring	(proof sketched above)
Isotypical of finite length	Matrix ring over division ring	(proof sketched above)
Semisimple of finite length	Direct sum of matrix rings over division rings	(proof sketched above)

The good thing about Artinian rings, as we shall soon see, is that *all* finitely generated modules over Artinian rings have finite length, and in particular, all semisimple finitely generated modules are semisimple of finite length, and thus the above can be applied.[12] Before plunging into the theory of Artinian rings, we shall take a bit of time to review what we have seen so far, and study examples of rings and modules that come up in the "real" world.

7. Jacobson radical, top and Nakayama's lemma

7.1. **Nakayama's lemma.** This is perhaps one of the most useful results in both commutative and noncommutative algebra, so it is useful to dwell on this result.

The "radicals" of a ring (commutative or noncommutative), such as the nilradical and the Jacobson radical, are collections of "small" elements, or elements which are close to zero in some sense. Thus, taking the quotient of a ring by a radical should yield a new ring which has fewer of the pathological 0-type elements, and in general more of invertibility.

Suppose I were a nilpotent ideal of a ring R and M were a R-module. Then IM would be a submodule, and clearly $IM = M$ is not possible, because we have $I^n M = 0$. Thus IM must be strictly smaller than M.

We can do this for nilpotent ideals, which are *really small*, but can we do this for the Jacobson radical? That could in general be a whole lot bigger. Is it true that if J is the Jacobson radical of R, and M is a left R-module, then $JM = M$ is impossible?

The answer is *yes, modulo finiteness assumptions.* However, we need not place the finiteness assumptions on R; we can place them on M relative to R.

Theorem 4 (Nakayama's lemma)**.** Let M be a finitely generated left R-module and J be the Jacobson radical of R. Then $JM = M$ implies that $M = 0$.

Proof. The idea is to show that in some sense JM is a *smaller* module than M; the precise sense in which we do this is to look at the size of a generating set.

Suppose $JM = M$. Then let n be the minimum possible size of a generating set for M. If $n = 0$ we have $M = 0$, and we are done. Otherwise, $n > 0$. Let $m_1, m_2, \ldots, m_n$ be a generating set for M of

[12] Actually, a semisimple finitely generated module over any ring has finite length; a proof of this is to be found in part 2.

size n. Then since $JM = M$, we can write m_1 as a linear combination of the m_is with coefficients in J, specifically:

$$m_1 = \sum_i a_i m_i$$

We can rewrite this as:

$$(1 - a_1)m_1 = \sum_{i=2}^{n} a_i m_i$$

But now since $a_1 \in J$, $1 - a_1$ is invertible, so we can multiply both sides by the inverse of $1 - a_1$ (on the left) and thus express m_1 as a linear combination of the remaining elements. Thus $m_2, m_3, \ldots, m_n$ form a generating set for M contradicting the minimality of n. □

The "cardinality" argument here may seem slick but what's really going on is that finitely generated modules, we can always find an *irredundant* generating set: a generating set from which no member can be thrown out. This is related to the basic fact that any descending chain of finite sets stabilizes at a finite stage, viz., an "Artinianness" holds for finite sets.[13]

7.2. **The Jacobson radical and the top.** We begin with a definition:

Definition (Top of a ring)**.** The **top of a ring**(defined) is defined as the quotient of the ring by its Jacobson radical.

The top of any ring has trivial Jacobson radical. This is a *top* if one thinks of invertible elements as in some sense being higher, and closer to 1, and the elements close to 0 as living inside the dreary depths.

The first thing we would like to do is study rings which are tops of themselves, viz., rings which have trivial Jacobson radical. These rings are sometimes termed semisimple (there are, however, alternative definitions of the notion of semisimplicity). After that, we shall try to see how the *whole* ring can be described from its top.

Just like for a ring, we can try looking at the top of a *module* over the ring, viz.,:

Definition (Top of a module)**.** Let R be a ring and M a R-module. Then the **top**(defined) of M is defined as the quotient module M/JM.

Note that by Nakayama's lemma, the top of any finitely generated module is nonzero, so there is some hope that the top of the module conveys nonzero information about the module.

Let's now state another corollary of Nakayama's lemma, which is almost as ubiquitous as the lemma itself:

Theorem 5 (Nakayama's lemma version 2)**.** Suppose R is a ring and M, N are finitely generated R-modules. Then if a R-homomorphism $f : M \to N$ induces an isomorphism from the top of M to the top of N, f is itself an isomorphism.

In other words, the property of being an isomorphism can be checked on the top.

8. Polynomial theory over noncommutative rings

8.1. **Free associative algebras over rings.** Given a ring R, the **free associative algebra**(defined) over R in variables $x_1, x_2, \ldots, x_n$, is defined as the ring whose elements are noncommutative polynomials in the variables x_i, with *coefficients* in R. We typically assume that elements of R commute with each of the x_i's, even if R is noncommutative. We shall denote the free associative algebra by $R < x_1, x_2, \ldots, x_n >$.

[13]This is often called "infinite descent", and is a contrapositive of sorts of the principle of mathematical induction

The free associative algebra in 1 variable over R, is the same as the polynomial ring in 1 variable. On the other hand, the free associative algebra in 2 variables over R is very different from the polynomial ring in 2 variables – in the free associative algebra x_1x_2 is linearly independent from x_2x_1.

Here are some cool facts about free associative algebras:

- The free associative algebra over any ring is free as a module over that ring, with a generating set being the set of all noncommutative monomials in the variables. By noncommutative monomial, we simply mean a string whose letters are from the variables.
- Let R be a commutative ring. Any R-algebra which, as an R-algebra is generated by n variables, is naturally a quotient of the free associative algebra on n variables, by a two-sided ideal. The natural map sends x_i to the i^{th} generator of the R-algebra.[14]
- Suppose the base ring is a field. Then any quotient of the free associative algebra has a basis over the field, of the images of noncommutative monomials. The idea is to start with the images of all monomials (which form a generating set for the quotient) and then go down to a linearly independent subset of that.

Here are some examples. Let R be a commutative ring (for now). The free associative algebra $R < x_1, x_2 >$ quotiented out by the ideal generated by $x_1x_2 - x_2x_1$, yields the polynomial ring $R[x_1, x_2]$.

8.2. **Standard monomials.** For convenience, let R be a commutative ring throughout.

Consider a quotient of the form $A = R < x_1, x_2, \ldots, x_n > /I$ where I is a two-sided ideal. The noncommutative monomials in the x_i's, viewed modulo I, generate A as a R-module. A first question would be: can we pass from noncommutative monomials to noncommutative monomials in a standard order.

Call a monomial a *standard monomial* if it is of the order $x_1^{a_1}x_2^{a_2}\ldots x_n^{a_n}$. One question is: under what conditions does the quotient have a generating set of standard monomials? Certainly, this is true for the polynomial ring, but it is true in a number of other cases, which we shall see. The rough idea is that in each case, we have a rule by which whenever we see x_jx_i with $j > i$, we can rewrite it as a linear combination of standard monomials.

8.3. **Some examples.** We list some examples for two variables, which we shall return to. For convenience, we call the two variables x and y.

Common name	Generator for two-sided ideal
Polynomial ring	$xy - yx$
Exterior algebra	$xy + yx$, x^2, y^2
Weyl algebra	$xy - yx - 1$
$A_1^-(R)$	$xy + yx - 1$
Translation ring	$xy - yx - x$

The polynomial algebra and exterior algebra admit direct generalizations to many variables; the Weyl algebra and translation ring need to be generalized with a little more care. Note that in each of these cases, yx can be expressed as a combination of standard monomials (1, x and xy) and hence the quotient admits a generating set of standard monomials. However, most of the other features of these rings are very different. We shall study these rings individually, and the relationships between them, in this section.

8.4. **Rings and derivations.** We define the notion of derivation, which will make a lot of what we are about to discuss easier:

Definition (Derivation)**.** Let R be a ring. A **derivation**(defined) of R is a map $d : R \to R$ with the property that d is an Abelian group map from R to R, and that $d(rs) = r(ds) + (dr)s$ (the Leibniz rule).

Derivations are commonly found in rings; for instance in the polynomial ring in two variables, differentiation in either variable is a "derivation". For noncommutative rings, there's an even more natural kind of derivation.

[14]Note it is crucial to observe that by our definition, elements of R must be in the center.

Definition (Inner derivation)**.** Let R be a ring. The **inner derivation**(defined) induced by an element $a \in R$ is defined as the map $x \mapsto ax - xa$.

As the terminology suggests, any inner derivation is a derivation. Inner derivations are special in many respects; one good thing is that two-sided ideals are closed under inner derivations. Here's a definition:

Definition (Lie ideal)**.** A subset of a ring is termed a **Lie ideal**(defined)[15] if it is a subgroup under addition, and is closed under all inner derivations.

When working with algebras over a base ring R, we call a subset a **Lie ideal**(defined) if it is a R-submodule of the algebra, and is closed under all inner derivations.

Inner derivations are to derivations the way inner automorphisms are to automorphisms. Here are some easy facts:

- Any two-sided ideal is a Lie ideal.
- A left ideal is a Lie ideal iff it is two-sided.[16]
- The center is a Lie ideal. More generally, any additive subgroup contained inside the center is a Lie ideal.
- A commutative ring has no nonzero inner derivations. Thus, the differentiation in either variable in a polynomial ring is an *outer derivation*.

The reason why we have set up all this terminology is to better understand what goes on in the examples we will look at. So now it's time to begin.

8.5. **The Weyl algebra.** In the next subsection, we shall see "concretely" how the Weyl algebra arises naturally as the algebra of "differential operators" on the polynomial ring; however, for now, we will play with it as an abstract ring. This shall also help us brush up abstract manipulation theory (or lack thereof).

Theorem 6. The Weyl algebra over a field of characteristic zero, has the property that any nonzero Lie ideal contains the element 1 of the base field (the base field will be denoted by $\mathbf{k}$).

Note that this will trivially imply that the Weyl algebra is simple, because it forces that any nonzero two-sided ideal contains 1 and is therefore the whole ring. However, the statement that we have made is *much stronger* than merely the statement of simplicity; for instance, matrix rings have a lot of nontrivial Lie ideals even though they are simple.

Instead of giving the entire proof, I'll list the main steps:

- By the fact that $yx = xy - 1$, we see that elements of the form $x^a y^b$ form a generating set for the Weyl algebra, as a $\mathbf{k}$-vector space.
- The inner derivation induced by x turns out to be "differentiation" with respect to y and the inner derivation induced by y turns out to be "differentiation" with respect to x (upto a sign change).
- If a Lie ideal contains a nonzero element, then repeated application of these inner derivations will yield an element of $\mathbf{k}$. This step uses characteristic zero: if the characteristic is *not* zero, then differentiating in either variable can send a nonzero "polynomial" to the zero polynomial.

This proof essentially also shows that the center of the Weyl algebra is precisely $\mathbf{k}$ when $\mathbf{k}$ has characteristic zero. When $\mathbf{k}$ has characteristic *not* zero, say characteristic p, then the center of the Weyl algebra is the polynomial ring generated by x^p and y^p. The Weyl algebra is a degree p^2 extension over it.

[15] The term "Lie ideal" arises from the fact that in the associated Lie algebra, we would get an ideal; however, this is largely irrelevant here

[16] This might be time, for the enterprising, to go back to the notion of "normal" left ideal and when such an ideal is forced to be two-sided

8.6. **The concrete realization of the Weyl algebra.** There are many ways of viewing the Weyl algebra as arising naturally from the polynomial ring. Here's one. The polynomial ring, in itself, is a well-behaved commutative ring, but it has a derivation, the so-called d/dx operator, and this derivation is not realized as an inner derivation in the polynomial ring. Ideally, we would like to embed the polynomial ring in a bigger ring where this derivation becomes inner.

If a derivation is already inner, it extends to something inner in any bigger ring, because there is already a witness for it in the smaller ring[17] However, an outer derivation in a smaller ring may become inner in an enlarged ring. This is analogous to the construction for groups is the "semidirect product".

For the polynomial ring $\mathbf{k}[x]$, we want to construct a bigger ring which contains $\mathbf{k}[x]$, in which differentiation with respect to x becomes inner derivation by some other element, say y. So let's stick in such a y artificially. It's easy to see that the derivation is completely determined by its effect on the element x and by the fact that it fixes the scalars, so the only "relation" we need to throw in is:

$$yx - xy = 1$$

This is the Weyl algebra, with the role of x and y reversed.

Thus, the Weyl algebra can be thought of as the "smallest" algebra containing the polynomial ring, in which differentiation with respect to x becomes an inner derivation.

There is an alternative way of saying what we said above, which is as follows:

Definition (Algebra of differential operators)**.** Let A be a R-algebra (R is a commutative ring). The **algebra of differential operators**(defined) of A is the R-subalgebra of $\mathrm{End}_R(A)$ (viz., R-module endomorphisms of A) generated by all the left multiplications by elements of A, and all derivations (inner or outer).

When A is a commutative R-algebra, the algebra of differential operators admits a natural "filtration" viz., we can talk of the "order" of a differential operator as the number of times one needs to do derivations to achieve it. In the particular case where $A = R[x]$, the algebra of differential operators that we get is the Weyl algebra.

8.7. **From another viewpoint.** Another viewpoint of what we've done above is that we have found a simple module for the Weyl algebra: namely the polynomial ring $\mathbf{k}[t]$, where y acts via left multiplication by t, and x acts via differentiation by t. This is a simple module, because repeated differentiation and left multiplication and addition can take us from any polynomial to any other polynomial.

Note that we've now made y act by multiplication and x by differentiation, because when we tried to do it the other way around, we got our signs reversed.

8.8. **Differential operators in many variables.** We proved that the Weyl algebra is precisely the algebra of differential operators on the polynomial ring in one variable; we can similarly compute the algebra of differential operators on the polynomial ring in many variables. The answer is that it is the algebra generated by elements of two kinds:

- Left multiplications in each of the variables
- Differentiation in each of the variables

If $\mathbf{k}[t_1, t_2, \ldots, t_n]$ is the polynomial algebra, y_i denotes left multiplication by t_i and x_i denotes differentiation in t_i, then clearly for $i \neq j$, x_i and x_j commute, y_i and y_j commute, x_i and y_j commute. For $i = j$, we have the relations $x_i y_i - y_i x_i = 1$.

8.9. **The exterior algebra.** Recall that the exterior algebra in two variables was given by the relations $xy + yx = 0$, $x^2 = 0$ and $y^2 = 0$. We can more generally talk of the exterior algebra in n variables, where the relations are given by $x_i^2 = 0$ and $x_i x_j + x_j x_i = 0$.

I will use R to denote the base ring and A to denote the exterior algebra.

I list some interesting observations:

- For the exterior algebra in n variables, the square of *every* element is 0, and *any* two elements anticommute. This should not be taken as immediate, because the conditions are specified only on basis elements, rather, it can be concluded from an explicit check.

[17] In fancy language, inner derivations are extensible

- All the generating elements $x_1, x_2, \ldots, x_n$ are *invariant* elements. In other words, $x_i A = A x_i$ where A denotes the whole algebra.
- The exterior algebra has order 2^n, and it has a basis of standard monomials with every monomial occurring to multiplicity at most 1.
- The exterior algebra is naturally a $\mathbb{Z}$-graded algebra where the gradation comes from the degree of a monomial.
- We can write:

$$A = A^{even} \oplus A^{odd}$$

  where the direct sum is as R-modules (R is the base ring). Elements of A^{even} are R-combinations of monomials of even degree, elements of A^{odd} are R-combinations of monomials of odd degree.
- Every element of A^{even} is in the center; every element of A^{odd} is an invariant element.
- All elements of A^{even}, as well as the monomial of length n, are in the center. For a field of characteristic not 2, the center of the exterior algebra is generated by these.
- Suppose $R = \mathbf{k}$, a field. As we had observed earlier, any invariant nilpotent is strongly nilpotent, and hence in the Jacobson radical. Since all the x_i's are invariant and nilpotent, they are all in the Jacobson radical; hence the Jacobson radical is precisely the ideal generated by all the x_i's.
- Thus, the exterior algebra over a field (of characteristic not 2) is a local ring with residue field equal to the field itself. Its center is also a local ring, with residue field equal to the same field.

8.10. **Differential operators on the exterior algebra.** We first need to compute what are the *derivations* of the exterior algebra. Again, because of the Leibniz rule, we only have freedom to determine where the x_i go, and any derivation must be an A-linear combination of "differentiation" in the x_is. It turns out that a little manipulation yields an algebra very similar to the Weyl algebra, but with signs reversed on all the relations.

Like the Weyl algebra, this is again simple, and its center is $\mathbf{k}$. However, since the exterior algebra, unlike the polynomial algebra, is finitely generated, this is a finite-dimensional simple ring over the field. The exterior algebra is a natural module of dimension 2^n for this ring, and a little bit of work yields that this is isomorphic to the matrix ring of order 2^n over $\mathbf{k}$.

8.11. **Relations that destroy symmetry.** Both the polynomial ring and the exterior algebra were nice in the following sense: any vector space automorphism on the vector space $\mathbf{k}x_1 \oplus \mathbf{k}x_2 \oplus \ldots \oplus \mathbf{k}x_n$ extended uniquely to an algebra automorphism. However, this symmetry is not held in general, *even in situations where there is symmetry in the variables*. For instance, in the ring:

$$\mathbf{k}[x_1, x_2]/(x_1^2, x_2^2)$$

Any ring automorphism must send x_1 either to a multiple of x_1 or to a multiple of x_2. This is closely related to the study of homotopy classes of self-maps of $S^2 \times S^2$; in fact the symmetric and exterior algebra are both special cases of algebras that arise in practice as graded-commutative rings in algebraic topology.

8.12. **The translation ring.** Let's conclude this half of the "flavour" by looking at a ring that has a different flavour from the rings discussed so far: the so-called *translation ring*, which is defined as $\mathbf{k}\langle x, y\rangle$ modulo the two-sided ideal generated by $xy - yx - x$.

We first note that x is an *invariant element*: the left ideal generated by x is the same as the right ideal generated by x. Moreover, x is clearly *not* invertible; it's also not a zero divisor. The two-sided ideal generated by $xy - yx - x$ lives inside the two-sided ideal generated by x, and quotienting out by that two-sided ideal yields the ring $\mathbf{k}[y]$.

Thus the translation ring has a proper nonzero two-sided ideal, namely the ideal generated by the image of x. However, something better is true for the translation ring: *any* nonzero two-sided ideal must contain a power of x. The proof of this again involves starting out with a nonzero element viewed as a standard monomial, and then repeatedly applying "derivations" to get to a polynomial that has no y-component; after that we apply derivations again till we get a power of x.

Index