Midterm 2 Review Notes

Christoph von Pezold, Russell Georgi

November 2, 2025

1 Generalised Kernels/Eigenspaces and the Jordan Normal Form

Definition.

Given a linear operator $\varphi : V \to V$, we have the following associated invariant subspaces (in increasing order of relevance):

1. The *eventual image*:
$$\text{evIm}\varphi = \bigcap_{n\in\mathbb{N}} \text{im}\varphi^n$$
2. The *generalised kernel*:
$$\text{gKer}\varphi := \bigcup_{n\in\mathbb{N}} \ker\left(\varphi^n\right)$$
3. For $\lambda \in \mathbb{C}^*$ the *generalised eigenspaces*:
$$V_\lambda := \text{gKer}\left(\varphi - \lambda I\right)$$
The λ for which $V_\lambda \neq 0$ are called the *generalised eigenvalues* of φ.

Note.

For V finite dimensional the nested sequence $\ker\varphi \subset \ker\varphi^2 \subset \cdots$ stabilises at $\ker\varphi^{\dim V}$ and as such in this setting the generalised kernel is just the kernel of $\varphi^{\dim V}$; similarly $\text{im}\varphi \supset \text{im}\phi^2 \supset \cdots$ stabilises at $\dim V$ and so the eventual image is just the image of $\varphi^{\dim V}$

Also note that every normal eigenspace is contained in a generalised eigenspace and every normal eigenvalue is a generalised eigenvalue, justifying the terminology.

Why do we care about these things? Essentially because they're our best hope for recovering a project that the existence of undiagonalisable linear operators kill the nicest form of: decomposing a vector space into nice invariant subspaces of a given linear operator so that we can nicely break up the action of that linear operator into more manageable chunks. This does not work in general for normal eigenspaces; although complex linear operators always have an eigenvalue, they need not decompose into a direct sum of eigenspaces. The crucial thing that prevents us from doing this in general is that if we try to start with our guaranteed eigenvalue and then decompose and induct, we run into the problem that the kernel and image of $\varphi - \lambda I$ can overlap; consider the map $\mathbb{R}^2 \to \mathbb{R}^2$ with matrix in the standard basis given by:
$$\mathcal{M}(\varphi) = \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}$$

The unique solution to the characteristic polynomial of this matrix is $\lambda = 1$, but $\ker(\varphi - I) = \text{span}\begin{pmatrix}1\\0\end{pmatrix} = \text{im}(\varphi - I)$ and so we very much do not have a direct sum decomposition. However! You in your problem set have proved that we always have a decomposition into the generalised kernel and the eventual image, which means we can prove the following theorem, for $\lambda_1, ..., \lambda_m$ the generalised eigenvalues of φ:

Theorem.

We have the following direct sum decomposition for any finite dimensional $\mathbb{C}$-vector space V and linear operator $\varphi : V \to V$:

$$V = \bigoplus_{i=1}^{n} V_{\lambda_i}$$

This is already very nice, and in fact reflects the two main things you will actually use generalised eigenspaces for. The first is that you have a direct sum decomposition all the time and as such can always make a basis of generalised eigenvectors, which is a convenient thing to be able to pull out of your hat and throw at a problem, especially because of the second nice thing which we saw used in this proof: that you can always make a generalised eigenvector into an eigenvector by applying a power of $\varphi - \lambda I$ to it, which gives you a way to get eigenvectors that you know a little bit more about. Of course the other reason we like them is because of the Jordan Normal Form, which we get to now:

Theorem.

Every linear operator φ on a $\mathbb{C}$-vector space has a basis in which it has a matrix that is a block diagonal composed of *Jordan Blocks*, those being blocks of the form:

$$\begin{pmatrix} \lambda & 1 & 0 & \cdots & 0 \\ 0 & \lambda & 1 & \cdots & 0 \\ 0 & 0 & \ddots & \ddots & \vdots \\ 0 & 0 & 0 & \lambda & 1 \\ 0 & 0 & 0 & 0 & \lambda \end{pmatrix}$$

Where each λ appearing in a jordan block is a generalised eigenvalue of φ

The proof of this proceeds basically by using quotients to analyze nilpotent operators, which the restriction to each generalised eigenspace of $\varphi - \lambda_i I$ are.

2 Bilinear Forms

Definition.

A *bilinear form* on a vector space V over k $B : V \times V \to k$ is a function linear in both inputs individually; that is for all $u, v, w \in V, a, b \in k$:

$$B(au + bv, w) = aB(u, w) + bB(v, w)$$
$$B(u, av + bw) = aB(u, v) + bB(u, w)$$

Note.

Bilinear forms are *not* linear functions on the vector space $V \times V = V \oplus V$! In particular, bilinearity implies that for a with $a^2 \neq a$ and $\mathbf{v} = (v, v) \in V \times V$ such that $B(\mathbf{v}) \neq 0$:

$$B(a\mathbf{v}) = B(av, av) = a^2 B(v, v) \neq aB(v, v)$$

This is a more general instance of a bilinear pairing $B : U \times V \to W$ which is defined entirely analogously; bilinear forms just happen to be a particularly useful special case. One thing that might help for thinking about them for the more matrix oriented among you is that if you pick a basis $(e_i)_{1 \leq i \leq n}$ of a vector space and associate to two vectors x, y the column vectors $X, Y \in k^n$ then letting $M_B = (a_{ij})$ with $a_{ij} = B(e_i, e_j)$ then the value of $B(x, y)$ can be brute computed to be:

$$B(v, w) = X^T M_B Y$$

Another[1] way of thinking about this is that we have a canonical association $\{\text{bilinear forms } V \times V \to k\} \to \text{Hom}(V, V^*)$:

Proposition.

We have a bijection[a]:

$$\begin{aligned} \{\text{bilinear forms } V \times V \to k\} &\to \text{Hom}(V, V^*) \\ B &\mapsto \varphi_B : v \mapsto B(v, \cdot) \end{aligned}$$

[a] actually an isomorphism, but you need tensors to see what that even means

The matrix of a bilinear form is then just the matrix of its associated linear map $V \to V^*$ with respect to the standard dual basis. These notions give us some ways of conceptualising some important properties bilinear forms can have:

Definition.

A bilinear form is

1. *nondegenerate* if φ_B is an isomorphism. This is equivalent to saying that for all $v \in V \setminus \{0\}$ there exists $w \in V$ with $B(v, w) \neq 0$. This corresponds to $\det M_B \neq 0$
2. *symmetric* if $B(v, w) = B(w, v)$ for all $v, w \in V$. This corresponds to M_B being a symmetric matrix i.e. $a_{ij} = a_{ji}$
3. *Skew-symmetric* or *alternating* if $B(v, w) = -B(w, v)$ for all $v, w \in V$; notice that this implies and in fact is equivalent to $B(v, v) = 0$ for all $v \in V$ – consider $B(v - w, v - w)$. This corresponds to M_B being a skew-symmetric matrix i.e. $a_{ij} = -a_{ji}$
4. *Positive-definite* if $B(v, v) \geq 0$ for all $v \in V$, with $B(v, v) > 0$ for all $v \in V \setminus \{0\}$. This does not have an easy reflection in the realm of matrices.

[1] better

Note.

Notice that φ_B feels much stronger than "for all $v \in V$ there exists $w \in V$ with $B(v,w) \neq 0$" and although they are equivalent, in practice it is often easier to think about proving the latter, but easier to use the former. So this is a good problem solving tactic with nondegenerate bilinear forms: prove that they are by using the 2nd definition, but use that they are using the first.

Using bilinear forms we can begin to define a couple concepts that make our linear algebra feel more similar to standard geometry, for example:

Definition.

Given a subspace $S \subset V$ and a symmetric or skew-symmetric bilinear form B the *orthogonal* of S with respect to B is:

$$S^{\perp} := \{w \in V \mid B(v,w) = 0 \ \forall v \in S\}$$

In order for this to work nicely though, we need to impose some restrictions on B. In particular, we would sure hope that the orthogonal doesn't intersect with the subspace itself, but this isn't even true in general for nondegenerate bilinear forms; consider $B : \mathbb{R}^2 \times \mathbb{R}^2 \to \mathbb{R}$ given by $B(X,Y) = \det(X,Y)$. To make things work nicely we need a more restrictive notion of taking the "product" of two vectors:

Definition.

An *inner product* $\langle \cdot, \cdot \rangle$ on a real vector space V is a positive definite symmetric bilinear form on V.

Note.

This is usually stated as positive-definite, nondegenerate, symmetric bilinear form[a], but this is actually redundant because positive-definite implies nondegenerate - just take v itself as your other vector!

[a] there's a great song about this on youtube somewhere

Exercise.

Axler 6A.2

Suppose $S \in \text{End}(V)$. Define $\langle ., . \rangle_1$ by

$$\langle u, v \rangle_1 = \langle Su, Sv \rangle$$

for all $u, v \in V$. Show that $\langle ., . \rangle_1$ is an inner product on V if and only if S is injective.

Definition.

The *norm* of a vector u in an inner product space V is:

$$\|u\| = \sqrt{\langle u, u \rangle}$$

There is also an analogous definition for complex vector spaces, although we have to give up full bilinearity because over $\mathbb{C}$ it prevents positive-definiteness. Instead:

Definition.

A *Hermitian Inner Product* H over a $\mathbb{C}$-vector space V is a conjugate-symmetric positive-definite sesquilinear[a] form on V. This means:

1. *conjugate symmetric*: $H(v,w) = \overline{H(w,v)}$ for all $v, w \in V$
2. *sesquilinear*: for all $\lambda \in \mathbb{C}, u, v, w \in V$:

$$H(\lambda v, w) = \bar{\lambda} H(v,w), H(v, \lambda w) = \lambda H(v,w), H(u, v+w) = H(u,v) + H(u,w)$$

3. *positive definite*: as before

We often write $H(\cdot,\cdot) = \langle \cdot, \cdot \rangle$ much like with standard inner products

[a] this is actually a bit redundant because the first implies the third but they're all important enough to state separately

Both of these kinds of inner product satisfy some very nice properties. Before we state all of them, it's worth knowing another bit of terminology. Whenever we introduce a new structure on vector spaces, it's good to know which kinds of bases respect that structure. The relevant construct here is:

Definition.

An *orthonormal basis* of an inner product space or hermitian inner product space V is a basis $(e_i)_{i \in I}$ such that:

$$\langle e_i, e_j \rangle = \delta_{ij} := \begin{cases} 1 & i = j \\ 0 & i \neq j \end{cases}$$

Note.

Any orthonormal set is linearly independent, so you actually only have to check orthonormality and spanning

We then have the following facts:

Proposition.

1. For any subspace $S \subset V$, $V = S \oplus S^{\perp}$
2. Every finite dimensional vector space admits an orthonormal basis
3. (Cauchy-Schwarz) $|\langle u, v \rangle| \leq \|u\|\|v\|$ for all $u, v \in V$ – with equality only when the two are colinear.

3 Operators on Inner Product Spaces

There are a couple kinds of special operators on inner product spaces that we care about. First:

Definition.

Given an inner product space[a] V and a linear operator $T : V \to V$, its *adjoint* $T^* : V \to V$ is the

unique linear operator such that for all $v, w \in V$:

$$\langle Tv, w\rangle = \langle v, Tw\rangle$$

[a] or Hermitian Inner Product Space

Note.

This can be seen as the conjugate of the transpose map $T^t : V^* \to V^*$ by the isomorphism of V with V^*. Over the reals its matrix is the transpose matrix, over $\mathbb{C}$ the conjugate of the transpose matrix.

Adjoints have some nice properties worth noting:

Proposition.

1. $\ker T^* = (\mathrm{im} T)^\perp$
2. $\mathrm{im} T^* = (\ker T)^\perp$
3. If λ is an eigenvalue of T then $\bar{\lambda}$ is an eigenvalue of T^*

Exercise.

Axler 7A.4

Suppose $T \in \mathrm{End}(V)$ and U is a subspace of V. Show that $U^\perp$ is invariant under T^* iff U is invariant under T.

We can use this to define a couple nice kinds of operators:

Definition.

1. An operator is *self-adjoint* if $T = T^*$; this corresponds to symmetric matrices in the real case.
2. An operator is *orthogonal* (in the real case) or *unitary* (in the complex case) if $T^* = T^{-1}$, that is $\langle Tu, Tv\rangle = \langle u, v\rangle$ for all $u, v \in V$.
3. A complex operator is *normal* if $TT^* = T^*T$; this contains the preceding cases

The main reason we care about self-adjoint operators is that, to put it lightly, they are *insanely nice*. In particular, they satisfy the following:

Theorem.

Spectral Theorem:

1. A real self adjoint operator is diagonalisable in an orthonormal basis; a complex self adjoint operator is diagonalisable in an orthonormal basis with real eigenvalues
2. An orthogonal real linear operator can be decomposed into 1 and 2-dimensional invariant subspaces on which it respectively acts by ± 1 and rotations. This is a consequence of complexification using:

3. A unitary complex linear operator is diagonalisable in an orthonormal basis with unit eigenvalues, that is $\lambda \in S^1$, that is $|\lambda| = 1$.

4. A normal complex linear operator is diagonalisable in an orthonormal basis.

Whenever you're not sure how to approach a problem involving inner product spaces, throw this at it and see what happens. With that in mind, it's worth noting:

Note.

T^*T is always self-adjoint

We can also carry out a more general classification of symmetric nondegenerate bilinear forms; these satisfy:

Proposition.

Every symmetric real bilinear form has a basis such that:

$$B(e_i, e_j) = \begin{cases} \pm 1 & i = j \\ 0 & i \neq j \end{cases}$$

In the complex case it's even better, we can get $B(e_i, e_j) = \delta_{ij}$

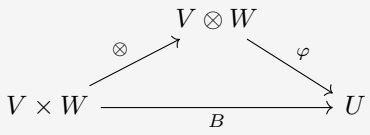
A similar classification can be done for skew-symmetric forms.

4 Tensors

Remember earlier when I said/wrote[2] that the correspondence between bilinear forms and linear maps $V \to V^*$ was just a bijection? I lied.[3] In fact it is an isomorphism, but in order to coherently phrase how we need to do a bit of work. Specifically, we want to define a way of talking about bilinear forms as a *linear* map from some vector space to k so that we can coherently talk about the space of all bilinear forms as a vector space that can be isomorphic to $\mathrm{Hom}(V, V^*)$. So essentially, we want a vector space that we can factor all of the bilinearity through and have the particular structure of the bilinear form be encoded by a linear map whereas the bilinearity part is encoded in the vector space structure. This motivates the following definition of the *tensor product*:

Definition.

Given vector spaces V, W the *tensor product* $V \otimes W$ is the unique[a] vector space with a bilinear map $\otimes : V \times W \to V \otimes W$ such that for any bilinear map $B : V \times W \to U$ there exists a linear map $\varphi : V \otimes W \to U$ making the following diagram commute:



[2] depending on what medium you're consuming this through
[3] unless you read the footnote

[a] up to isomorphism

This, for fancy category theory reasons involving limits of diagrams and the existence of objects satisfying universal properties (aka abstract nonsense[4]), is actually a valid definition, and once you are used to tensors in fact the best one. It is easier to see it in the following way however:

Note.

The tensor product $V \otimes W$ is, given a basis (v_i) of V and (w_j) of W, the space with basis the formal symbols $v_i \otimes w_j$ and relations enforced so that $\otimes$ is linear. It is *no more than this* in the sense that only those relations are satisfied and no more. Another way of saying this is that $V \otimes W$ is the quotient of the absolutely enormous space with basis the formal symbols $v \otimes w$ for $v \in V, w \in W$ by relations enforcing bilinearity, namely $(av) \otimes w - v \otimes (aw), (u+v) \otimes w - u \otimes w - v \otimes w$ etc.

Note.

We call elements of the form $v \otimes w \in V \otimes W$ *pure tensors*, but these are NOT all of them – the space is the space of all finite sums of pure tensors, as otherwise it wouldn't be a vector space. The minimal number of pure tensors needed to express a given tensor is called its rank.

We can then always write bilinear maps as $B(v, w) = \varphi(v \otimes w)$ for some linear map $\varphi : V \otimes W \to U$. The following proposition is vitally important:

Proposition.

$$V^* \otimes W \cong \mathrm{Hom}(V, W)$$

Proof Sketch: The isomorphism is given by $\ell \otimes w \mapsto \ell(\cdot)w$ and extending by linearity.

We also have a couple intuitive isomorphisms:

Proposition.

1. $V \otimes W \cong W \otimes V$
2. $(V \otimes W)^* \cong V^* \otimes W^*$
3. $(U \oplus V) \otimes W \cong (U \otimes W) \oplus (V \otimes W)$

It is in this setting that we can say what we mean by the statement "the space of bilinear forms is isomorphic to the space of linear maps $V \to V^*$." The space of bilinear forms is just the space of *linear* forms on $V \otimes V$, i.e. $\mathrm{Hom}(V \otimes V, k)$. The statement is then:

Proposition.

$$\mathrm{Hom}(V \otimes V, k) \cong \mathrm{Hom}(V, V^*)$$

Proof: Use the map given earlier. Alternatively, use the properties of tensors just above; they are

[4] which I use as a term of endearment but you are free to see otherwise

sufficient since $\text{Hom}(V \otimes V, k) = (V \otimes V)^*$.

Note.

This is a special case of a very, *very* cool fact called *tensor-hom adjunction*. You ***do not need to know this***, but googling it is worth it.

You also proved on a problem set that:

Proposition.

The trace of a matrix corresponds to the coordinate free construction of the linear map

$$\begin{aligned} \kappa : V^* \otimes V &\to k \\ \ell \otimes v &\mapsto \ell(v) \end{aligned}$$

under the isomorphism $\text{Hom}(V, W) \xrightarrow{\sim} V^* \otimes W$

This reflects a general way in which tensors allows us to define the standard constructions of linear algebra in a coordinate free[5] way. Another way to see this is with the *determinant*, but that goes into multilinear algebra which isn't on the midterm so we won't be going over it here.

5 Classifying Abelian Groups

You studied Modules over a Ring to some extent; you won't need to know most of that theory for the midterm, but one piece of information is relevant: Abelian Group correspond exactly to $\mathbb{Z}$-modules.[6] To see this, write the group operation of G as addition, and define for $n \in \mathbb{Z}$ and $g \in G$ with $0 := e \in G$ the identity element and $-g \in G$ the inverse element of g:

$$ng := \begin{cases} \overbrace{g + g + \cdots + g}^{n \text{ times}} & n > 0 \\ 0 & n = 0 \\ \underbrace{(-g) + (-g) + \cdots + (-g)}_{n \text{ times}} & n < 0 \end{cases}$$

We can thus use the theory of $\mathbb{Z}$-modules to understand abelian groups. A key tool in doing so is the following proposition:

Proposition.

Every finitely generated $\mathbb{Z}$-module M with generating set $\{e_1, ..., e_n\}$ is a quotient of $\mathbb{Z}^n$ by first isomorphism theorem applied to the surjective homomorphism:

$$\varphi : (a_1, ..., a_n) \mapsto \sum_{i=1}^{n} a_i e_i$$

[5] hence better

[6] In a very, very strong sense: the characterisation I'm about to give is an isomorphism of categories, which is such a stupidly strong notion of being the same that category theorists basically never use it; mostly they use something called equivalence of categories instead

Moreover $\ker \varphi$ is free meaning there exists $T : \mathbb{Z}^m \to \mathbb{Z}^n$ such that:

$$M \cong {}^{\mathbb{Z}^n}\!/_{\mathrm{im}T}$$

This fact is useful in proving the following very powerful theorem about finitely generated abelian groups, which is a special case of a more general concept from algebra called the *Structure Theorem*, so we often call this the Structure Theorem for Finitely Generated Abelian Groups:

Theorem.

For any finitely generated abelian group G there exist integers $r, n_1, n_2, ..., n_k$ such that:

$$G = \mathbb{Z}^r \oplus \bigoplus_{i=1}^{k} {}^{\mathbb{Z}}\!/_{n_i\mathbb{Z}}$$

If one moreover requires that $n_1 \mid n_2 \mid \cdots \mid n_k$ or that the n_i are all prime powers, this representation is unique.

Note.

We use the direct sum symbol to emphasise that abelian groups are $\mathbb{Z}$-modules, but you can write the direct product symbol as well and be just as correct; so we could also write:

$$G = \mathbb{Z}^r \times {}^{\mathbb{Z}}\!/_{n_1\mathbb{Z}} \times \cdots \times {}^{\mathbb{Z}}\!/_{n_k\mathbb{Z}}$$

6 Group Actions and Conjugation

Before we get to the proverbial meat and potatoes of this part of the course, it's worth just noting down a generally important definition:

Definition.

1. Given a group G, the *commutator subgroup* $[G, G]$ the subgroup generated by *commutators* $[a, b]$ which are defined as $[a, b] = aba^{-1}b^{-1}$. This is always a normal subgroup.[a]
2. The *abelianisation* of a group G is the quotient ${}^{G}\!/_{[G,G]}$

[a] If $x \in [G, G]$ and $g \in G$ then $gxg^{-1} = gxg^{-1}x^{-1}x = [g, x]x \in [G, G]$

Note.

Notice that a and b commute if and only if $[a, b] = e$, so quotienting out by the commutator is effectively forcing every pair of elements to commute, hence why the abelianisation is called that.

Now on to the main part. The key here is the following definition:

Definition.

An *action* of a group G on a set S is a homomorphism $\rho : G \to \text{Perm}(S)$ where $\text{Perm}(S)$ is the group of permutations of S, isomorphic to the symmetric group on $|S|$ letters. This is a succinct way of writing that we have a map:

$$\begin{aligned} G \times S &\to S \\ (g, s) &\mapsto g(s) = g \cdot s \end{aligned}$$

With the properties that:

1. $\forall s \in S : e(s) = s$
2. $\forall g, h \in G, s \in S : h(g \cdot s) = (hg) \cdot s$

We sometimes write that $G \curvearrowright S$ in this setting.

Note.

Action is exactly the right word for this concept. You want to think about each group element being assigned a way of acting on the set in such a way that composing these actions is the same as first composing in the group.

Example.

The prototypical examples of actions of a group action are:

1. $S_n \curvearrowright \{1, ..., n\}$. That S_n acts on $\{1, ..., n\}$ is almost definitional to S_n and $\text{Perm}(\{1, ..., n\})$ is exactly the definition of S_n so we can simply take our homomorphism ρ to be identity.
2. D_n or more generally any group of geometric symmetries acts on the geometric objects of which it is the symmetry group
3. $\mathbb{Z}/n\mathbb{Z}$ acts on the vertices of an n sided polygon by rotation
4. The real line $\mathbb{R}$ acts on itself by translation, but also on S^1 by rotation.

Some important definitions about actions:

Definition.

1. An action is *faithful* if ρ is injective; this means that every group element acts genuinely differently from every other, meaning the action faithfully captures the structure of the group, hence the name.
2. An action is *transitive* if for all $s, t \in S$, there is a $g \in G$ such that $g(s) = t$; you can reach any element of the set from a different element of the set via the action.
3. The *orbit* $O_s = \text{Orb}(s)$ of an element $s \in S$ with respect to an action $G \curvearrowright S$ is $\{gs : g \in G\}$, the set of all elements one can reach from s via G. Transitive actions only have one orbit. Being in the same orbit is an equivalence relation, hence we have a partition $S = \bigsqcup_{s \in S} O_s$
4. The *stabiliser* $\text{Stab}(s)$ of an element $s \in S$ with respect to an action $G \curvearrowright S$ is $\{g \in G : g(s) = s\}$.

This is a subgroup of G.

Note.

We call actions "faithful" because often the reason we want to study actions is because they actually allow us to study the structure of the group itself; this will be especially true in a very cool subject you get to a bit after the midterm called *representation theory*, where you consider not just actions by permutations but actions by *linear maps*. The point being that actions can give us information about the group itself, which is why we want the action to be faithful. Nonfaithful actions still tell us something though, it's just that it's not *really* G that they're telling us things about, but rather $G/\ker \rho$, a group which has an induced *faithful* action on S. This is basically just first isomorphism theorem again.[a]

[a] you'd be surprised just how much of algebra is philosophically the same thing as first isomorphism theorem

A quick relevant fact:

Proposition.

$\mathrm{Stab}(g \cdot s) = g\mathrm{Stab}(s)g^{-1}$, elements of the same orbit have conjugate stabilisers.

Understanding the stabilisers of elements is very important in a lot of applications so knowing that since conjugation gives a group isomorphism we only really need to understand the stabilisers for representative elements of each orbit is very powerful. Here's another way we interrelate orbits and stabilisers, the aptly named orbit-stabiliser theorem:

Theorem.

We have for any $s \in S$ a well-defined, G-equivariant bijection:

$$\begin{aligned} \varepsilon : G/\mathrm{Stab}(s) &\xrightarrow{\sim} O_s \\ [a] &\mapsto a(s) \end{aligned}$$

Where G-equivariant means that $g(\varepsilon([a])) = \varepsilon([ga])$

Corollary.

$|G| = |\mathrm{Stab}(s)||O_s|$

Note.

All of the maps in that theorem are well defined only because we're quotienting with respect to the *stabiliser* in particular; this ensures that $[ga]$ is well defined, because cosets only differ by things that do nothing to s.

Another useful fact for using actions to study symmetry is the following combinatorial identity, known as Burnside's Lemma. We'll need a quick definition first:

Definition.

For an action $G \curvearrowright S$, the set of *fixed points* of an element $g \in G$ is:

$$S^g := \{s \in S \mid g(s) = s\}$$

Note.

More generally if $H \leq G$ is a subgroup then S^H is the intersection of S^h for $h \in H$; only the former is necessary to know though.

We can now state the lemma:

Lemma.

For an action $G \curvearrowright S$, we have that the number of orbits is given by the following formula:

$$\#(\text{Orbits}) = \frac{1}{|G|} \sum_{g \in G} |S^g|$$

Note.

Usually the way we use this in practice is that we want to understand how many different possible objects there are up to some symmetry condition, so we make the set of symmetries that we want to consider the things "the same" under into a group and then determine the number of orbits of that group – the number of essentially different objects – via Burnside's lemma, because usually the number of fixed points for each element is easier to calculate, especially because conjugate elements always have the same number of fixed points.

Now we know how groups can act on sets. You know what else is a set? A group[7]! So how might a group act on itself? There's one obvious candidate:

Proposition.

A group acts on itself faithfully by left multiplication, $\rho(g)(h) = gh$.

Corollary.

Every finite group is isomorphic to a subgroup of S_n for some $n \in \mathbb{Z}^+$.

This action is however not very useful other than that fun fact, which is called *Cayley's Theorem*.[8] A far more useful action, because it fixes identity which is generally a nice property, is by *conjugation*:

Proposition.

Every finite group acts on itself by conjugation, $\rho(g)(h) = ghg^{-1}$

The orbits and stabilisers of these actions have special names:

Definition.

1. The *conjugacy class* $C_g = \text{ccl}(g)$ is defined to be the orbit of g under the conjugation action, that is the set of elements conjugate to g.

[7] strictly speaking it's a pair of a group and an operation or if you want to be really abstract a one point category, but just apply the forgetful functor

[8] Or, if you're a category theorist, a very very special case of Yoneda's Lemma. That's a fun google, but again like most of these footnotes not too relevant

2. The *Centraliser* $Z(g)$ is the stabiliser of g under the conjugation action, equivalently it is the set of elements of G that commute with g; notice that the center $Z(G) = \bigcap_{g \in G} Z(g)$.

Note.

Observe that the elements in the center $Z(G)$ that commute with everything are exactly the members of the singleton conjugacy classes

From these definitions and our previously established facts about actions, the following propositions follow:

Proposition.

For every $g \in G$, $|C_g| \mid |G|$ since $|C_g| = |G|/|Z(g)|$.

Proposition.

We have the following equality, known as the *class equation* of G:

$$|G| = \sum_{\text{Conjugacy Classes } C} |C|$$

These two propositions together tell us a lot about a finite group based on its order. For example:

Proposition.

Any group of order a power of a prime p has $|Z(G)| \geq p$

Proof Sketch Otherwise since $|Z(G)|$ is the number of singleton conjugacy classes the class equation would reduce to something nonzero modulo p

Corollary.

Any group of order p^2 for prime p is abelian.

A more substantive example from lecture classifies groups of order 8:

Proposition.

The groups of order 8 are as follows:

$$\mathbb{Z}/8\mathbb{Z}, \mathbb{Z}/4\mathbb{Z} \times \mathbb{Z}/2\mathbb{Z}, \left(\mathbb{Z}/2\mathbb{Z}\right)^3, D_4, Q$$

For Q the quaternion group $\{\pm 1, \pm i, \pm j, \pm k\}$ with $i^2 = j^2 = k^2 = -1$.

Before we get back to the purely group theory stuff, one last note about groups acting on things other than themselves:

Proposition.

If $G \subset \mathrm{SO}(3)$ is finite, that is G is a finite group of rotations of $\mathbb{R}^3$, then by considering its action on its poles of rotation we retrieve that G is isomorphic to one of:

$$\mathbb{Z}/n\mathbb{Z}, D_n, A_4, S_4, S_5$$

Last couple things you need to know are about symmetric groups, which since they're the prototypical groups that act on things should come as no surprise:

Proposition.

1. S_n is generated by transpositions (in fact *adjacent* transpositions of the form $(m\ m+1)$)
2. Every $\sigma \in S_n$ has a unique up to reordering expression as a product of disjoint cycles.
3. $\sigma, \tau \in S_n$ are in the same conjugacy class if and only if their cycle lengths are equal[a]

[a] because $\sigma(a_0\ a_1\ \cdots\ a_k)\sigma^{-1} = (\sigma(a_0)\ \sigma(a_1)\ \cdots\ \sigma(a_k))$

Definition.

If the disjoint cycle decomposition of σ is given by a product of a k_1 cycle with a k_2 cycle and so on up to a k_ℓ cycle, then the *cycle lengths* of σ are $(k_1, k_2, \ldots, k_\ell)$.

Every symmetric group has associated to it a very important subgroup, induced by the homomorphism $\text{sgn} : S_n \to {}^{\mathbb{Z}}\!/_{2\mathbb{Z}}$ that counts the number of transpositions modulo 2:

Definition.

The *alternating group* A_n is defined as follows:

$$A_n := \ker(\text{sgn})$$

Some facts about A_n:

Proposition.

A conjugacy class of S_n consisting of even permutations is either 1 or 2 conjugacy classes in A_n, determined by whether $Z(\sigma) \subset A_n$ – in which case it splits into two conjugacy classes, otherwise it remains 1 – which is equivalent to all cycle lengths being odd and distinct

Theorem.

For $n \geq 5$ A_n is simple.