Lie Algebras: summary note

Kevin McGerty

Trinity 2023

We briefly review the main ideas and results of the course.

i) The definition of a Lie algebra (crucially the Jacobi identity) and examples $\mathfrak{gl}_V$, $\mathfrak{sl}_V$ *etc.*. Basic structures such as ideals, quotients and isomorphism theorems.

ii) Nilpotent and Solvable Lie algebras: Engel's theorem and Lie's theorem. The main technique to prove results for such algebras is induction (*e.g.* on dimension using a short exact sequence in the inductive step, or by using the lower central series/derived series). For Lie's theorem we need to work over an algebraically closed field k of characteristic zero, and the assumption that k is of this form is in place for essentially all the remaining results in the course. This is for two (different and equally important) reasons:

   1. If a field k is algebraically closed, then any endomorphism of a k-vector space has an eigenvalue. (This is actually an "if and only if" statement).
   2. If k has characteristic zero, and $\lambda \in \mathsf{k}$, then for any integer $n \in \mathbb{Z}_{>0}$ we have $n.\lambda = 0$ if and only if $\lambda = 0$. This comes up because we can often calculate the trace of a linear endomorphism α by choosing an appropriate basis, and then if we know α has a single eigenvalue λ (where k is algebraically closed), we see that if $\mathrm{tr}(\alpha) = \dim(V).\lambda = 0$ if and only if $\lambda = 0$. (This is, for example, how the proof of Lie's Lemma works.)

iii) It's also worth noting that the proof of Lie's Lemma uses another idea – the "trace trick" – which comes up in many proofs in the course (it's one of the principal technical tools in the course): If $x, y, z \in \mathfrak{g}$ are elements of the Lie algebra $\mathfrak{g}$, and (V, ρ) is a representation of $\mathfrak{g}$, then if $z = [x, y]$ and $\rho(z), \rho(x), \rho(y)$ all preserve a subspace W of V, the trace of $\rho(z)$ on W is zero, since it is the commutator $[\rho(x), \rho(y)]$. Proofs which use this trick combine it with some other way of computing $\mathrm{tr}_W(\rho(z))$, which is in most cases via a knowledge of the generalised eigenspace decomposition of W with respect to $\rho(z)$. Note it is crucial that $\rho(x)$ and $\rho(y)$ preserve W not just $\rho(z)$: consider the representation of $\mathfrak{sl}_2(\mathsf{k})$ on k^2 with standard basis $\{e_1, e_2\}$. If $\{e, h, f\}$ is the standard basis of $\mathfrak{sl}_2$, then h preserves $\mathsf{k}.e_1$, and $h = [e, f]$, but f does not preserve W and $\mathrm{tr}_W(h) = 1$.

iv) Representations of nilpotent Lie algebras: linear endomorphisms decompose a vector space naturally into generalised eigenspaces. Representations of nilpotent Lie algebras behave in a similar way: every representation (V, ρ) of a nilpotent Lie algebra $\mathfrak{h}$ decomposes into a direct sum of subrepresentations V_α where V_α has only one composition factor, the one-dimensional representation k_α ($\alpha \in (\mathfrak{h}/D(\mathfrak{h}))^* = D(\mathfrak{h})^0$). V_α may thus also be characterised as the simultaneous (*i.e.* the intersection of the) generalised $\alpha(h)$-eigenspaces for the elements of $\rho(\mathfrak{h})$. That is, $V_\alpha = \{v \in V : (\rho(h) - \alpha(h))^n(v) = 0, \forall\, n \geq \dim(V), \forall\, h \in \mathfrak{h}\}$.

v) Cartan subalgebras: these subalgebras allow us to use the representation theory of nilpotent Lie algebras to study the structure of general (finite-dimensional) Lie algebras. The proof of the existence of Cartan subalgebras uses the notion of a regular element, and Engel's theorem.

The motivation for the definition is the following: we have seen that the representation theory of nilpotent Lie algebras is, in a sense, easy – representations are a direct sum of their generalised weight spaces. Since any 1-dimensional subspace of a Lie algebra $\mathfrak{g}$ is a subalgebra isomorphic to $\mathfrak{gl}_1$, if $\mathfrak{g}$ is nonzero, it contains nontrivial nilpotent subalgebras. If $\mathfrak{h}$ is any such subalgebra, by restriction, any $\mathfrak{g}$-representation (V, ρ) decomposes as an $\mathfrak{h}$-representation into $V = \bigoplus_{\lambda \in \Psi_V} V_\lambda$ where $\Psi_V \subseteq D(\mathfrak{h})^0 \subseteq \mathfrak{g}^*$ are the weights of $\mathfrak{h}$ which occur as composition factors of V and V_λ is the λ-*generalised weight space* of V, that is, the isotypical subrepresentation of V associated to the isomorphism class of irreducible representation given by λ. In particular, taking $(V, \rho) = (\mathfrak{g}, \text{ad})$, the adjoint representation, gives the decomposition $\mathfrak{g} = \bigoplus_{\lambda \in \Psi_\mathfrak{g}} \mathfrak{g}_\lambda$. Since $\mathfrak{h}$ is nilpotent, $\mathfrak{h} \subseteq \mathfrak{g}_0$, and so $\Psi_\mathfrak{g} = \{0\} \cup \Phi$, that is, Φ denotes the *roots*, *i.e.*, the (possibly empty) set non-zero weights of $\mathfrak{h}$ that occur in $\mathfrak{g}$.

Since, for any $\mathfrak{h}$-representations V and W and weights $\lambda, \mu \in D(\mathfrak{h})^0$ we have $V_\lambda \otimes W_\mu \subseteq (V \otimes W)_{\lambda+\mu}$ and if $\theta\colon V \to W$ is a homomorphism of $\mathfrak{h}$-representations, then $\theta(V_\lambda) \subseteq W_\lambda$, it is easy to see that $[\mathfrak{g}_\alpha, \mathfrak{g}_\beta] \subseteq \mathfrak{g}_{\alpha+\beta}$ and $\mathfrak{g}_\alpha(V_\mu) \subseteq V_{\alpha+\mu}$, so that $\mathfrak{g}$ and its representations are “graded” by the weights of $\mathfrak{h}$.

Indeed if $\mathfrak{h} \subseteq \mathfrak{h}' \subseteq \mathfrak{g}$ where $\mathfrak{h}'$ is again nilpotent, then the decompositions obtained using $\mathfrak{h}'$ will be finer than those obtained using $\mathfrak{h}$, thus we obtain more information about the structure of $\mathfrak{g}$ by using nilpotent subalgebras which are maximal with respect to containment. However, the decomposition $V = \bigoplus_{\lambda \in \Psi_V} V_\lambda$ actually gives a decomposition of V into $\mathfrak{g}_0$-subrepresentations ($\mathfrak{g}_0$ is a subalgebra because $[\mathfrak{g}_0, \mathfrak{g}_0] \subseteq \mathfrak{g}_{0+0} = \mathfrak{g}_0$). In general $\mathfrak{h} \subseteq \mathfrak{g}_0$ but the containment will usually be strict. In particular, $\mathfrak{g}_0$ need not be a nilpotent Lie algebra, and so to understand the structure of $\mathfrak{g}$ using the decomposition $\mathfrak{g} = \mathfrak{g}_0 \bigoplus_{\alpha \in \Phi} \mathfrak{g}_\alpha$ we would need to understand the structure of $\mathfrak{g}_0$, at least enough to understand its representations.[1]

If we can guarantee that $\mathfrak{h} = \mathfrak{g}_0$, then it is easy to see that this implies that $\mathfrak{h}$ is a maximal nilpotent subalgebra with respect to containment, and clearly in this case $\mathfrak{g}_0$ is nilpotent. The condition that $\mathfrak{h}$ is self-normalizing is in fact equivalent (for nilpotent subalgebras) to the condition that $\mathfrak{g}_0 = \mathfrak{h}$, that is, that $\mathfrak{h}$ is a Cartan subalgebra. Indeed in the proof of the existence of Cartan subalgebras one considers the 0-generalised eigenspaces $\mathfrak{g}_{0,x}$ of $\text{ad}(x)$ for $x \in \mathfrak{g}$. This is equivalent to considering the the abelian subalgerba $\mathfrak{h}_x = \mathsf{k}.x$ of $\mathfrak{g}$ and taking its 0-generalised weight space $\mathfrak{g}_0$. The existence theorem is proved by showing two things: for any x, the subalgebra $\mathfrak{g}_{0,x}$ is self-normalizing, and that provided x is *regular*, so that $\dim(\mathfrak{g}_{0,x})$ is minimal, then $\mathfrak{g}_{0,x}$ is nilpotent, and hence a Cartan subalgebra.

vi) Cartan’s Criterion for solvable and semisimple Lie algebras: Bilinear forms play a critical role in the course through the Killing form. The proof of Cartan’s criterion (or rather Lemma 5.2.10) relies on the “trace trick” where the vanishing of the trace of a commutator shows that the restriction of a weight λ to the subspace $[\mathfrak{g}_\alpha, \mathfrak{g}_{-\alpha}] \subseteq \mathfrak{h}$ must equal a rational multiple of the root α. To see how this works, let $V = \bigoplus_{\lambda \in \Psi_V} V_\lambda$ be the decomposition of V as an $\mathfrak{h}$-representation, and set $W = \bigoplus_{-p \leq k \leq q} V_{\lambda + k\alpha} \subseteq V$ where $\lambda - (p+1)\alpha \notin \Psi_V$, and $\lambda + (q+1)\alpha \notin \Psi_V$, so that W is stable under the action of $\mathfrak{g}_\alpha \oplus \mathfrak{g}_{-\alpha}$. Then if $z \in [\mathfrak{g}_\alpha, \mathfrak{g}_{-\alpha}]$, the trace of $\rho(z)$ must vanish on W since $\rho(z) \in [\rho(\mathfrak{g}_\alpha), \rho(\mathfrak{g}_{-\alpha})]$ must be a sum of commutators. On the other hand, for any $h \in \mathfrak{h}$, we have $\text{tr}_{V_\mu}(\rho(h)) = \dim(V_\mu).\mu(h)$, since $\rho(h)$ has $\mu(h)$ as its only eigenvalue on V_λ. We thus see that the vanishing of $\text{tr}_W(\rho(z))$ gives an equation relating $\lambda(z)$ and $\alpha(z)$ which readily establishes that upon restricting to $[\mathfrak{g}_\alpha, \mathfrak{g}_{-\alpha}]$ λ becomes equal to a rational multiple of α. This is the same strategy is used later when studying α-root strings through β for α, β roots of a semisimple Lie algebra.

[1] As an example, if $\mathfrak{g} = \mathfrak{b}_n$ denotes the upper-triangular $n \times n$ matrices and $\mathfrak{h} = \mathfrak{n}_n$ the strictly upper-triangular matrices, then $\mathfrak{h}$ is a nilpotent subalgebra which is maximal with respect to containment, but it is also an ideal, hence $\mathfrak{g} = \mathfrak{g}_0$, and so the decomposition of $\mathfrak{g}$ one obtains by viewing it as an $\mathfrak{h}$-representation is trivial. Indeed if $\mathfrak{h}$ is any nilpotent ideal of $\mathfrak{g}$ then $\mathfrak{g} = \mathfrak{g}_0$ as an $\mathfrak{h}$-representation.

vii) The Jordan decomposition: This gives a generalisation of the Jordan decomposition of a linear map to the case of any semisimple Lie algebra. The idea of the proof that the decomposition exists is to show that any derivation of a semisimple Lie algebra is inner, and that the semisimple and nilpotent parts of a derivation are again derivations (which is true for any Lie algebra). The fact that any derivation of a semisimple Lie algebra is inner is equivalent to showing that you cannot extend $\mathfrak{gl}_1$ by a semisimple Lie algebra in a non-trivial way, and this follows from Cartan's criterion in the same way that one shows that a semisimple Lie algebra is a direct sum of simply (nonabelian) Lie algebras.

viii) The compatibility of the Killing form and the Cartan decomposition of a semisimple Lie algebra is the key to the detailed structure of semisimple Lie algebras. We use the trace trick of *iii*) above a number of times to get information about the root spaces.

ix) Weyl's theorem: the semisimplicity of finite dimensional representations of semisimple Lie algebras follows roughly the pattern of the standard proof of the corresponding theorem for representations of finite groups: using Hom-spaces the problem of finding a complementary subrepresentation for an arbitrary subrepresentation $U \leq V$ is reduced to the problem of finding a complementary subrepresentation to the subrepresentation of invariants $V^{\mathfrak{g}}$. There is a natural candidate[2] – the subrepresentation $\mathfrak{g}.V$ – and by using the Casimir operators we can prove that $V = V^{\mathfrak{g}} \oplus \mathfrak{g}.V$ for any finite-dimensional representation of a semisimple Lie algebra $\mathfrak{g}$. This decomposition yields a natural $\mathfrak{g}$-homomorphism from V to $V^{\mathfrak{g}}$, which is precisely what the "averaging" operator $|G|^{-1}\sum_{g\in G}\rho(g)$ yields in the case of finite group representations, so it can be thought of as the analogue of that operator.

x) Root systems: this part of the course has a different feel from the rest as it uses just elementary facts about reflections in Euclidean vector spaces (*i.e.* a vector spaces with a positive definite bilinear form, and hence a notion of angle and distance). The existence of a base and the analysis of the Weyl group are the crucial results, *e.g.* that the Weyl group acts transitively on the collection of all bases in a root system. Thess results shows that root systems are determined up to isomorphism by the Cartan matrix.

xi) The classification theorems: Putting everything together, semisimple Lie algebras are classified by Cartan matrices (or Dynkin diagrams) and these can be completely classified. We don't have time in the course to prove all of this, but the statements(only!) which are required should be known, and there are two main results:

1. The uniqueness result: the procedure which associates a root system to a semisimple Lie algebra is well-defined, that is, the root system attached to a semsimple Lie algebra is unique. The procedure involves a choice of a Cartan subalgebra, so it is enough to know that if any two Cartan subalgebras are conjugate by an automorphism of the semisimple Lie algebra. This is in fact true for any finite dimensional Lie algebra over an algebraically closed field of characteristic zero (but uses some ideas we did not cover in the course).
2. The existence result: Given any root system, there is a semisimple Lie algebra with that root system. This can be proved in a number of ways and only requires the use of techniques which we used in the course.

[2] $V/\mathfrak{g}.V$ is the largest quotient of V on which $\mathfrak{g}$ acts trivially. If the trivial representation occurs in V only in $V^{\mathfrak{g}}$, as it must if V is completely reducible, then $V^{\mathfrak{g}}$ will map isomorphically onto $V/\mathfrak{g}.V$, so that $V = V^{\mathfrak{g}} \oplus \mathfrak{g}.V$.