REPRESENTATIONS OF NILPOTENT LIE ALGEBRAS

KEVIN MCGERTY

1

This note gives an account of the representation theory of nilpotent Lie algebras following the presentation used in the lecture videos which was slightly different to that used in the online lecture notes.

Lie's theorem shows that if $\mathfrak{g}$ is a nilpotent Lie algebra, its irreducible representations are all one-dimensional. For any Lie algebra, the one-dimensional representations are given by $(\mathfrak{g}/D\mathfrak{g})^*$, and when $\mathfrak{g}$ is nilpotent we will call these *weights* of $\mathfrak{g}$.

Definition 1.1. If $\mathfrak{g}$ is a Lie algebra and $\lambda \in (\mathfrak{g}/D\mathfrak{g})^*$ then we will write k_λ for the one dimensional representation given by $\lambda\colon \mathfrak{g} \to \mathfrak{gl}_1$.

By the Jordan-Holder theorem, any finite dimensional representation V is built from irreducible representations in the sense that you can find an increasing sequence $0 = F_0 < F_1 < \ldots F_k = V$ of subrepresentations such that F_i/F_{i-1} is irreducible, and the multiplicity of occurence of the irreducible representation is independent of the choice of filtration. For solvable Lie algebras, this along with Lie's theorem, is as complete a description of a general representation as we will be able to obtain: in more concrete terms it say that if V is a representation of a solvable Lie algebra $\mathfrak{g}$, then there is a basis of V with respect to which the matrices of the linear maps $\rho(x)$ are all upper triangular (since all irreducible representations of $\mathfrak{g}$ are one-dimensional). The irreducibles which occur in a composition series for V are called the *weights of* V. For nilpotent Lie algebras however, we will be able to get more information:

Definition 1.2. For any weight λ of a nilpotent Lie algebra $\mathfrak{g}$ and any $\mathfrak{g}$-representation (V, ρ), we set

$$V_\lambda = \{v \in V : (\rho(x) - \lambda(x))^k(v) = 0, \text{for some } k \in \mathbb{Z}_{>0}\}$$

Thus V_λ is the intersection of the generalized eigenspaces of $\rho(x)$ with eigenvalue $\lambda(x)$ over all $x \in \mathfrak{g}$.

Remark 1.3. We can also rephrase the definition of V_λ more representation-theoretically: by Lie's theorem, we know that all the irreducible representations of $\mathfrak{g}$ are one-dimensional, and the only one-dimensional representation which can occur in the space V_λ is (up to isomorphism) k_λ, since for each $x \in \mathfrak{g}$ the linear map $\rho(x)$ can only have the eigenvalue $\lambda(x)$ on any composition factor[1] of V_λ. Thus V_λ is the largest subrepresentation of V whose composition factors are all isomorphic to k_λ. In particular, if we write $[V : \mathsf{k}_\lambda]$ for the multiplicity with which k_λ occurs in a composition series of V, then

$$\dim(V_\lambda) \leq [V : \mathsf{k}_\lambda]. \tag{1.1}$$

The key to our description of representations of a nilpotent Lie algebra is the following:

Lemma 1.4. *Let (V, ρ) be a representation of a nilpotent Lie algebra $\mathfrak{g}$ and let $x \in \mathfrak{g}$. The if $V = \bigoplus_\lambda V_{\lambda,x}$ is the generalized eigenspace decomposition of V with respect to $\rho(x)$, each $V_{\lambda(x)}$ is a subrepresentation of V.*

Proof. Fix $\lambda \in \mathsf{k}$. To see that $V_{\lambda,x}$ is a $\mathfrak{g}$-subrepresentation we the following equation, established in Section 7 of the main lecture notes: Writing $\rho\colon \mathfrak{g} \to \mathfrak{gl}(V)$ for the action of $\mathfrak{g}$ on V, we have

$$(\rho(x) - \lambda)^k \rho(y)(v) = \sum_{r=0}^{k} \binom{k}{r} \rho(\mathrm{ad}(x)^r(y))(\rho(x) - \lambda)^{k-r}(v), \quad \forall x, y \in \mathfrak{g}, v \in V.$$

Now since $\mathfrak{g}$ is nilpotent $\mathrm{ad}(x)$ is nilpotent on $\mathfrak{g}$ and hence there is some N for which $\mathrm{ad}(x)^N = 0$. Similarly if we take $v \in V_{\lambda(x)}$ there is some M for which $(\rho(x) - \lambda)^M(v) = 0$. Thus if $k \geq N + M$ either $\mathrm{ad}(x)^r = 0$ or $(\rho(x) - \lambda)^{k-r}(v) = 0$, thus each term on the right-hand side of the above equation vanishes, and hence $\rho(y)(v) \in V_\lambda$ as required. □

[1] If $\phi\colon W \to W$ is a linear map with λ as its only eigenvalue, $\phi - \lambda$ is nilpotent on W, and hence on U_1/U_2 for any subspaces $U_1 > U_2$ of W preserved by ϕ.

Theorem 1.5. *Let $\mathfrak{g}$ be a nilpotent Lie algebra and suppose that (V,ρ) is a $\mathfrak{g}$-representation of $\mathfrak{g}$. Then*
$$V = \bigoplus_{\lambda\in(\mathfrak{g}/D\mathfrak{g})^*} V_\lambda.$$

The proof in the main online lecture notes is different from the one given in the lecture videos. We give an account of that second proof here, along with what is essentially the same proof as in the main lecture notes, but expressed in a more representation theoretic way:

First proof:

We use by induction on $\dim(V)$, the case $\dim(V) = 1$ being trivial. Thus suppose we have V with $\dim(V) > 1$. Then by Lie's theorem, we may find a one-dimensional subrepresentation L of V. Now if we set $W = V/L$ then $\dim(W) = \dim(V) - 1$, so by induction we know that $W = \bigoplus_{\lambda\in\Psi} W_\lambda$, where Ψ is the finite set of elements of $(\mathfrak{g}/D\mathfrak{g})^*$ for which $W_\lambda \neq 0$. Set $\tilde{W}_\lambda = q^{-1}(W_\lambda)$ where $q\colon V \to W$ is the quotient map. Now $(L, \rho_{|L})$ is isomorphic to k_μ for some $\mu \in (\mathfrak{g}/D\mathfrak{g})^*$, hence we obtain a short exact sequence:
$$0 \longrightarrow \mathsf{k}_\mu \xrightarrow{\;i\;} \tilde{W}_\lambda \xrightarrow{\;q\;} W_\lambda \longrightarrow 0$$
where i is induced by the isomorphism $\mathsf{k}_\lambda \cong L$ and q, by abuse of notation, is the restriction of q to $\tilde{W}_\lambda$. We claim it is enough to show that if $\lambda \neq \mu$ then this sequence splits, *i.e* we can find an $s\colon W_\lambda \to \tilde{W}_\lambda$ such that $q\circ s_\lambda = \mathrm{id}_{W_\lambda}$. Indeed if this is the case then $s_\lambda(W_\lambda)\oplus L$ decomposes $\tilde{W}_\lambda$ as a direct sum of subrepresentations. There are now two cases: either $\mu \notin \Psi$, in which case $V = L \oplus \bigoplus_{\lambda\in\{Psi} s_\lambda(W_\lambda)$, or $\mu \in \Psi$ in which case $V = \tilde{W}_\mu \oplus \bigoplus s_\lambda(W_\lambda)$. Now clearly each of the spaces L, $\tilde{W}_\mu$ and $s_\lambda(W_\lambda)$ have only one composition factor, (k_μ for the first two and k_λ for the third, so they lie in V_μ and V_λ respectively. But by considering, for each of the cases, a composition series for V compatible with the direct sum decomposition, it is clear that for each $\eta \in \Psi\cup\{\mu\}$ the multiplicity $[V : \mathsf{k}_\eta]$ is equal to the dimension of the corresponding summand, hence (see Remark 1.3) the summands are in fact equal to the subspaces V_λ as required.

It thus remains to show that the short exact sequences split. We do this in the next Lemma.

Lemma 1.6. *Let $\mathfrak{g}$ be a nilpotent Lie algebra, and suppose that we have a short exact sequence of $\mathfrak{g}$ representations*
$$0 \longrightarrow \mathsf{k}_\mu \xrightarrow{\;i\;} V \xrightarrow{\;q\;} W \longrightarrow 0$$
where W has only one composition factor, k_λ, and $\lambda \neq \mu$. Then the sequence splits, that is, there is a homomorphism of $\mathfrak{g}$-representations $s\colon W \to V$ with $q \circ s = id_W$.

Proof. Since $\lambda \neq \mu$, we may pick $x \in \mathfrak{g}$ with $\lambda(x) \neq \mu(x)$. Now on W, the action of x has unique eigenvalue $\lambda(x)$ while on $i(\mathsf{k}_\mu)$ it has eigenvalue $\mu(x)$. It follows that if $V = V_{\lambda(x)} \oplus V_{\mu(x)}$ is the generalized eigenspace decomposition of V with respect to the action of x, then $V_{\mu(x)} = i(\mathsf{k}_\mu)$, so since $\mathrm{im}(i) = \ker(q)$, the map q is injective on $V_{\lambda(x)}$ and hence gives an isomorphism $q_{|V_{\lambda(x)}}\colon V_{\lambda(x)} \to W$. Its inverse $s\colon W \to V_{\lambda(x)}$ will give a splitting map provided $V_{\lambda(x)}$ is a subrepresentation[2]: Indeed as $q_{|V_{\lambda(x)}}\colon V_{\lambda(x)} \to W$ is a homomorphism of $\mathfrak{g}$-representations, its inverse s is automatically a $\mathfrak{g}$-homomorphism, since if q is invertible and for all $x \in \mathfrak{g}$ we have $q\circ x = x\circ q$ then pre- and post- composing both sides with $s = q^{-1}$ we see that $(s\circ q)\circ(x\circ s) = (s\circ x)\circ(q\circ s)$ and hence $x\circ s = s\circ x$ as required.

□

Second proof:

Write Ψ for the set of one-dimensional representations of $\mathfrak{g}$ which occur a composition series for V. Then we may find $x \in \mathfrak{g}$ with $\lambda(x) \neq \mu(x)$ whenever $\lambda \neq \mu$: Indeed the set $\Delta(\Psi) = \{\lambda - \mu : \lambda,\mu \in \Psi, \lambda \neq \mu\}$ is finite and thus $\bigcup_{\eta\in\Delta(\Psi)}\{x \in V : \eta(x) = 0\}$ is a finite union of hyperplanes (*i.e.* subspaces of codimension 1) in $\mathfrak{g}/D(\mathfrak{g})$ and hence[3] it cannot be all of $\mathfrak{g}/D(\mathfrak{g})$. Any $x \in \mathfrak{g}$ whose coset $x + D\mathfrak{g}$ lies in the complement of this union of hyperplanes will suffice.

Now let $V = \bigoplus_{\eta\in\mathsf{k}} V_{\eta,x}$ be the generalized eigenspace decomposition of V with respect to the action of $\rho(x)$. Now by taking a composition series of V as a $\mathfrak{g}$-representation[4] we see the eigenvalues of $\rho(x)$ are exactly $\{\lambda(x) : \lambda \in \Psi\}$, and clearly $V_\lambda \subseteq V_{\lambda(x)}$, so that since the generalized eigenspaces form a direct sum so must the V_λ. It remains to show that $V_\lambda = V_{\lambda(x)}$. But $V_{\lambda(x)}$ is a $\mathfrak{g}$-representation, and if k_μ is a

[2]This is a general fact, a short exact sequence $0 \longrightarrow U \xrightarrow{\;i\;} V \xrightarrow{\;q\;} W \longrightarrow 0$ of representations of any Lie algebra splits if and only if V decomposes as $V = i(U)\oplus W'$, with the splitting map being given by the inverse of the restriction of q to W' – see the appendix of the main online lecture notes on short exact sequences.

[3]This is true for a vector space over an infinite field, and any algebraically closed field in infinite.

[4]As each irreducible is one-dimensional, this is a complete flag, so taking a basis of V compatible with the flag the matrices associated to $\rho(\mathfrak{g})$ are upper triangular with the diagonal entries giving the action on the corresponding composition factor.

composition factor of $V_{\lambda(x)}$ then $\mu(x)$ is an eigenvalue of $\rho(x)$ on $V_{\lambda(x)}$, so that $\mu(x) = \lambda(x)$ and hence by our choice of x we must have $\mu = \lambda$. Thus every composition factor of $V_{\lambda(x)}$ is isomorphic to k_λ and so $V_{\lambda(x)} \subseteq V_\lambda$ and since we already observed $V_\lambda \subseteq V_{\lambda(x)}$ we are done.

For completeness, we include a proof of the elementary result about hyperplanes used in the previous proof.

Lemma 1.7. *Suppose that* k *is an infinite field and* V *is a* k*-vector space. The if* $\{H_i : 1 \leq i \leq n\}$ *is a finite collection of hyperplanes in* V *we have* $\bigcup_{i=1}^n H_i \subsetneq V$.

Proof. Prove this by induction on the number of hyperplanes (the case $n = 1$ being trivial).

Let $\{H_i : 1 \leq i \leq n\}$ be a collection of n hyperplanes in V. Now if $H_1 \subseteq H_k$ for any $k > 1$, then $\bigcup_{i=1}^n H_i = \bigcup_{i=2}^n H_i$, and so we are done by induction. Thus we may assume that $H_1 \cap H_k$ is a hyperplane in H_1 and so again by induction $H_1 \supsetneq \bigcup_{i=2}^n (H_1 \cap H_i)$. It follows that there is a vector $v_1 \in H_1 \backslash (\bigcup_{i=2}^n H_i)$. Now consider $w_t = v_1 + tw$ where we pick $w \notin H_1$ (which we can do by the $n = 1$ case). Now suppose that there exist $s, t \in \mathsf{k}$ with $s \neq t$ but $w_s, w_t \in H_k$ for some k. Then $w_s - w_t = (s-t).w \in H_k$ and so $w \in H_k$, and hence $k \geq 2$. But then $v_1 = w_s - s.w \in H_k$, which contradicts our choice of v_1. Moreover, since $w \notin H_1$, it is clear that $w_t \in H_1$ if and only if $t = 0$. Thus for each $k \geq 1$, H_k contains at most one of the vectors on the line $L = \{v_1 + tw : t \in \mathsf{k}\}$, hence as k is infinite, infinitely many points on the line L lie in none of the hyperplanes H_k as required. □