Representations of real reductive groups.

Kevin Buzzard

April 26, 2012

Last modified March 2007.

What are these notes?

I am sick of not knowing what words like "square-integrable", "principal series" and "discrete series" mean when applied to real reductive Lie groups. So I'm ploughing through Knapp's "Representation theory of semisimple groups" and here is a summary of what I learnt so far. Warning: these notes get vaguer/sloppier as they go on.

1 Reductive Lie algebras over R.

Let $\mathfrak{g}$ be a real Lie algebra with the property that $\mathfrak{g}_{\mathbf{C}} := \mathfrak{g} \otimes_{\mathbf{R}} \mathbf{C}$ is reductive in the usual sense of complex Lie algebras. You can rig it so that $\mathfrak{g}$ is a subalgebra of a matrix algebra over the complexes (or even over the reals!) and furthermore that this subalgebra is stable under the map $X \mapsto \theta X := -\overline{X}^t$: the point is that you can rig it so that the Cartan involution becomes identified with this involution. This shouldn't be logically necessary but the convention becomes very useful when it comes down to examples.

Now over the complexes we know the story well; all the Cartan subalgebras are conjugate and if you choose one then you get roots, the root spaces are 1-dimensional and, when direct summed with the Cartan, give the whole Lie algebra. Moreover after another choice, namely an ordering of the roots, we get positive and simple roots, and we get a lattice of weights, and the dominant ones (for our choice of simple roots) give us representations of the Lie algebra $\mathfrak{g}_{\mathbf{C}}$, and we get all the irreducible ones this way.

Over the reals things are slightly different, especially if you want the analogous decomposition to take place over $\mathbf{R}$.

The first new invariant is what this Cartan involution is doing on $\mathfrak{g}$ itself. Embed $\mathfrak{g}$ as a subalgebra of a real matrix algebra and break it up into eigenspaces $\mathfrak{k} \oplus \mathfrak{p}$ (the Cartan decomposition), where $\mathfrak{k}$ (this is a gothic k) is the subspace where $\theta = 1$ and $\mathfrak{p}$ has $\theta = -1$. It turns out that $\mathfrak{k}$ is the Lie algebra of a maximal compact subgroup K of G, and G is diffeomorphic to $K \times \mathfrak{p}$ (also the Cartan decomposition). In particular G and K have the same cohomology.

Note however that if G and H are algebraic groups over $\mathbf{R}$, and are forms of each other, then their Lie algebras (which of course have the same dimension and the same complexification) might very well have very different Cartain decompositions!

Examples: if $G = \mathrm{SL}_3(\mathbf{R})$ then $\mathfrak{g}$ is trace zero 3x3 real matrices, $\mathfrak{k}$ is the antisymmetric ones (3-dimensional), and $\mathfrak{p}$ is the symmetric ones (5-dimensional). If $G = \mathrm{SU}(2,1)$ formed with the form which is 1s up the antidiagonal then $\mathfrak{k}$ is 4-dimensional (the maximal compact is $S(U(2) \times U(1))$ which is isomorphic to $U(2)$) as is $\mathfrak{p}$. If $G = \mathrm{SU}(3)$ defined as $g \in \mathrm{GL}_3(\mathbf{C})$ with $g\overline{g}^t = 1$ then $\mathfrak{g} = \mathfrak{k}$ and $\mathfrak{p} = 0$.

Because θ "commutes with the Lie bracket", we see that $[\mathfrak{k}, \mathfrak{k}] \subseteq \mathfrak{k}$ and $[\mathfrak{p}, \mathfrak{p}] \subseteq \mathfrak{k}$ too, and $[\mathfrak{k}, \mathfrak{p}] \subseteq \mathfrak{p}$. Now we can get a new real Lie algebra! Think of it as $\mathfrak{k} \oplus i\mathfrak{p}$ in $\mathfrak{g}_{\mathbf{C}}$. In other words change the sign of $[p, q]$ if $p, q \in \mathfrak{p}$. This new Lie algebra is all fixed under θ so corresponds to a compact connected Lie group! It is somehow "the compact form of $\mathfrak{g}$".

Example: split tori change to non-split tori, $U(1)$ stays as it is, $\mathrm{SL}_n(\mathbf{R})$ becomes $SU(n)$ as does $SU(a,b)$ for $a+b=n$, $\mathrm{SL}_n(\mathbf{C})$ is just a form of $\mathrm{SL}_n(\mathbf{R}) \times \mathrm{SL}_n(\mathbf{R})$ so it will become $SU(n)^2$, and so on. Any compact group stays the same.

Weil's unitary trick says that the finite-dimensional smooth complex representations of a connected real reductive Lie group G are "the same" as the reps of its compact form, because they are both the same as complex reps of the complexified Lie algebra of G (which is the same as that of its compact form).

One checks that the irreducible unitary representations of a compact connected reductive Lie group K are all finite-dimensional and are parametrised by dominant integral weights in the weight lattice associated to the complexified Lie algebra $\mathfrak{k}_{\mathbf{C}}$ (this is not quite right because of finite centre issues, really you get a lattice which is commensurable with the dominant integral weights I guess).

Later on I'll tell you the Borel-Weil theorem, which gives an explicit construction of all the representations of K! It's Theorem 5.29 of Knapp and it's really important because it's what lies at the heart of holomorphic discrete series representations. However first I have to explain how root systems work for real Lie algebras, and how to get around the problem that for a random Cartan subgroup, the eigenspaces may not be defined over the reals (for example $SO_2(\mathbf{R})$ is in $\mathrm{SL}_2(\mathbf{R})$ but one can't now decompose the Lie algebra into eigenspaces for $\mathrm{so}_2(\mathbf{R})$ because the adjoint action isn't diagonalisable over $\mathbf{R}$; the three eigenvalues over $\mathbf{C}$ are 0 and $\pm 2i$).

2 Restricted root systems.

For Lie algebras over the complexes, a Cartan subalgebra is a maximal abelian subalgebra and in the reductive case all such things are conjugate and hence isomorphic and in particular have the same dimension. In the real case all Cartan subalgebras have the same dimension (tensor up to $\mathbf{C}$!) but they are not in general conjugate over $\mathbf{R}$. For example consider the Lie algebras of the split (diagonal) torus in $\mathrm{SL}_2(\mathbf{R})$ and the non-split torus $\mathrm{SO}_2(\mathbf{R})$; these can't be conjugate because they are both 1-dimensional but the eigenvalues for the split torus on the adjoint representation are real and those for the non-split one are pure imaginary.

More generally, the game here is that we have a real reductive Lie algebra $\mathfrak{g} = \mathfrak{k} \oplus \mathfrak{p}$ and we want to try and do root systems without tensoring up to $\mathbf{C}$, and we are going to have to be careful with our choice of Cartan. What one does first is chooses a maximal abelian subspace $\mathfrak{a}_p$ of $\mathfrak{p}$ (I don't know what that p is doing there, but there will be an $\mathfrak{a}$ later which is a subspace of $\mathfrak{a}_p$). Even though this is probably not a Cartan, it's what we'll use. What is happening here is that on the Lie group side we're choosing a maximal *split* torus.

While we're here, here are some basic definitions: the *rank* of G is the real dimension of a Cartan, and the *real rank* of G is the dimension of $\mathfrak{a}_p$. Note of course that the real rank is at most the rank, and for $G = \mathrm{SL}_n(\mathbf{R})$ the rank and real rank coincide. If G is the complex points of an algebraic group over $\mathbf{C}$ then the rank is twice the real rank, because $\mathfrak{a} \oplus i\mathfrak{a}$ is a Cartan. Appendix C of Knapp "rep theory of semisimple Lie groups" contains a table of ranks and real ranks.

It turns out that the action of all of $\mathfrak{a}_p$ on the adjoint representation will be diagonalisable over the reals (we're talking about $\mathfrak{a}_p$ acting on $\mathfrak{g}$ here) and so $\mathfrak{g}$ decomposes as a direct sum of $\mathfrak{g}_\lambda$ for λ running through the real vector space dual of $\mathfrak{a}_p$.

Note that $\mathfrak{g}_0$ (which is of course the centralizer of $\mathfrak{a}_p$ in $\mathfrak{g}$) contains $\mathfrak{a}_p$ but in this situation there is no reason to expect that $\mathfrak{g}_0 = \mathfrak{a}_p$ because there could be a subspace of $\mathfrak{k}$ which centralizes $\mathfrak{a}_p$. We will see this in the $SU(2,1)$ case below. In fact there's a name for the error: let $\mathfrak{m}_p$ be the centralizer of $\mathfrak{a}_p$ in $\mathfrak{k}$; then by definition $\mathfrak{g}_0 = \mathfrak{m}_p \oplus \mathfrak{a}_p$.

Why not bump up $\mathfrak{a}_p$ by adding elements of $\mathfrak{k}$? Well you can envisage $\mathfrak{k}$ giving problems because the eigenvalues of elements of $\mathfrak{k}$ on the adjoint representation will probably be pure imaginary and we want decompositions over the reals here.

Any λ with $\lambda \neq 0$ and $\mathfrak{g}_\lambda \neq 0$ is called a restricted root, and $\mathfrak{g}_\lambda$ for these λ is a restricted root space. Note again that there is no reason to expect that the restricted root spaces over the reals are 1-dimensional, and we'll see an example of this below.

Note that if one complexifies $\mathfrak{a}_p$ then it is still abelian, but there is no reason to expect that it is maximal abelian in $\mathfrak{g}_{\mathbf{C}}$, so it's not a Cartan subalgebra. However it will be contained within a Cartan subalgebra and one can do the usual decomposition of $\mathfrak{g}_{\mathbf{C}}$ into root spaces. The inclusion from the complexification of $\mathfrak{a}_p$ into the Cartan induces a surjection on dual spaces and it turns out that to understand the restricted roots you should think of them as restrictions of roots!

Note that the resulting root system might not be reduced! We'll see this below in the $SU(2,1)$ case, which is perhaps the simplest example of this phenomenon?

Examples: I worked out explicit examples for the forms of $\mathrm{SL}_3(\mathbf{R})$.

1) If $\mathfrak{g} = \mathrm{sl}_3(\mathbf{R})$ then as I said above, $\mathfrak{g}$ is the trace zero 3x3 real matrices, $\mathfrak{k}$ is the antisymmetric ones and $\mathfrak{p}$ is the symmetric ones. It turns out that one can let $\mathfrak{a}_p$ be the diagonal matrices in $\mathfrak{p}$, which is 2-dimensional, and the restricted roots are just the same as the roots, and all the restricted root spaces are 1-dimensional over the reals, and $\mathfrak{g}_0 = \mathfrak{a}_p$, so $\mathfrak{m}_p = 0$.

2) If $\mathfrak{g} = \mathrm{su}(3)$ then $\mathfrak{g} = \mathfrak{k}$ and so $\mathfrak{p} = 0$ and so $\mathfrak{a}_p = 0$ and so there are no restricted roots and $\mathfrak{g}_0 = \mathfrak{g} = \mathfrak{m}_p$. Note that $\mathfrak{g}_0$ is not even abelian here, the rank is 2 but the real rank is 0. Note also that "M_p is compact but not abelian" (see later for M_p, which is just $SU(3)$ in this case).

3) If $\mathfrak{g} = \mathrm{su}(2,1)$ then you can work it all with matrices. I used the unitary group associated to J, the matrix with 1s up the antidiagonal. Then a typical element of $\mathfrak{g}$ is

$$\begin{pmatrix} a & b & ix \\ d & \overline{a} - a & -\overline{b} \\ iy & -\overline{d} & -\overline{a} \end{pmatrix}$$

with a, b, d complex and x, y real, and a typical element of $\mathfrak{k}$ is

$$\begin{pmatrix} ix & b & iy \\ -\overline{b} & -2ix & -\overline{b} \\ iy & b & ix \end{pmatrix}$$

and a typical element of $\mathfrak{p}$ is

$$\begin{pmatrix} x & b & iy \\ \overline{b} & 0 & -\overline{b} \\ -iy & -b & -x \end{pmatrix}$$

and one can let $\mathfrak{a}_p$ be the diagonal elements of $\mathfrak{p}$, because this is clearly abelian and furthermore one checks easily that $\mathrm{diag}(1,0,-1)$ doesn't commute with anything else in $\mathfrak{p}$, and now one finds that $\mathfrak{g}_0$ is 2-dimensional, containing all the diagonal elements of $\mathfrak{g}$, $\mathfrak{m}_p$ contains $\mathrm{diag}(i,-2i,i)$, and the restricted root space has size 4 in the line and looks like $\{\pm\alpha, \pm 2\alpha\}$ and the root spaces corresponding to $\pm\alpha$ are also 2-dimensional. What is going on here is that the usual complex root system is the vertices of a hexagon in 2-space and it projects onto the root system of size 4 in 1-space via projection of a plane onto a well-chosen line. Note that the reason some of the restricted root spaces have got real dimension 2 is because they are coming from more than one complex root basically. Note also in this example that $\mathfrak{g}_0$ is 2-dimensional, one dimension of course coming from $\mathfrak{a}_p$ and the other coming from its centralizer in $\mathfrak{k}$.

Back to the general case. Note that one can put an ordering on the restricted roots, giving us a set of positive restricted roots, and hence a real Lie subalgebra $\mathfrak{n}_p$, and the Iwasawa decomposition of $\mathfrak{g}$ is $\mathfrak{g} = \mathfrak{k} \oplus \mathfrak{a}_p \oplus \mathfrak{n}_p$. In fact this lifts to the level of groups, with $G = K \times A \times N$ a diffeomorphism.

3 Parabolics.

Now let G be semisimple. We have enough machinery to define the standard *minimal parabolic* of G: we choose $\mathfrak{a}_p$ in $\mathfrak{g}$, exponentiate it to get $A \subseteq G$ (note that $A_p \cong (\mathbf{R}^n_{>0})$ for n the real rank), we let M_p be the centralizer of A_p in K, we choose an ordering on the restricted roots and we let N_p denote the subgroup coming from the subalgebra $\mathfrak{n}_p$ associated to the positive roots. Then $M_p A_p N_p$ is the standard minimal parabolic, and a minimal parabolic is something conjugate to this.

A standard parabolic is any subgroup of G containing the standard minimal parabolic; a parabolic is a subgroup of G containing a minimal parabolic. Any parabolic S has a decomposition $S = MAN$ with A a split torus, N unipotent and M reductive with compact centre but possibly not connected, and S is diffeomorphic with $M \times A \times N$.

It turns out that one can classify the standard parabolics: if we consider the restricted root system and choose an order and take the simple roots, then the parabolics biject with the subsets of the simple roots. The way it works is that if Π_S is a subset of the simple rootthen for $\lambda \in \Pi$ we have $\lambda \in \Pi_S$ iff $\mathfrak{g}_{-\lambda} \subseteq \mathfrak{m}$, the Lie algebra of M, and (I'm not entirely sure what this means) $\mathfrak{a}$ is the orthocomplement of the H_λ for $\lambda \in \Pi_S$, and if we consider $\mathfrak{a}$ acting on $\mathfrak{g}$ via the adjoint representation then $\mathfrak{n}$ is a direct sum of certain eigenspaces, and $\mathfrak{g} = \mathfrak{a} \oplus \mathfrak{m} \oplus \mathfrak{n} \oplus \theta\mathfrak{n}$, but be warned that we don't get an abstract root system in this setting because the axioms don't quite work; imagine a parabolic corresponding to $4 = 2 + 1 + 1$ in $\mathrm{SL}_4(\mathbf{R})$; then $\mathfrak{a}$ is 2-dimensional but the characters of $\mathfrak{a}$ are not really normalised correctly; if a general element of $\mathfrak{a}$ is $\mathrm{diag}(\lambda, \lambda, \mu, -2\lambda - \mu)$ then the roots are $\pm(1, -1)$, $\pm(3, 1)$ and $\pm(2, 2)$ which are topologically an abstract root system but in practice aren't one. Note also that the analytic Weyl group $N_K(\mathfrak{a})/Z_K(\mathfrak{a})$ only has order 2 and doesn't permute the Weyl chambers transtively.

4 The Borel-Weil theorem.

If K is a compact connected Lie group (note that this actually implies that its Lie algebra is reductive, as K is a closed subgroup of a unitary group by Peter-Weyl) then here is the representation theory of K. First take its Lie algebra $\mathfrak{k}$ over the reals. Now tensor up to $\mathbf{C}$ and call the resulting Lie algebra $\mathfrak{g}$. Now $\mathfrak{g}$ can be thought of as a reductive Lie algebra *over the reals* and if you do this then you can ask about its Cartan decomposition, which of course is $k \oplus ik$, but $\mathfrak{g}$ can also be thought of as a complex reductive Lie algebra and as such is the Lie algebra of a complex reductive algebraic group $G_{\mathbf{C}}$. The way we've rigged it, K is a maximal compact subgroup of $G_{\mathbf{C}}(\mathbf{C})$.

Examples: if $K = U(n)$ then $G_{\mathbf{C}} = \mathrm{GL}_n(\mathbf{C})$. If $K = SO(n)(\mathbf{R})$ then $G_{\mathbf{C}}$ is *not* anything like $\mathrm{SL}_n(\mathbf{R})$ (which isn't even the complex points of a complex algebraic group!), it's $SO(n)(\mathbf{C})$. And so on.

Now let $\mathfrak{a}$ be a maximal abelian subalgebra of $\mathfrak{p} = i\mathfrak{k}$ and extend to $\mathfrak{h} = \mathfrak{b} \oplus \mathfrak{a} = i\mathfrak{a} \oplus \mathfrak{a} = \mathfrak{m} \oplus \mathfrak{a}$. Choose an ordering; get $\mathfrak{n}$; set $\mathfrak{b} = \mathfrak{m} \oplus \mathfrak{a} \oplus \theta\mathfrak{n}$ (the lower triangular matrices) and exp up to $B \subseteq G$.

If λ is a dominant weight then λ gives us a 1-dimensional representation of AM and hence of B. Now consider

$$\Gamma(\lambda) := \{F : G_{\mathbf{C}}(\mathbf{C}) \to \mathbf{C} \text{ holomorphic } : F(xb) = \lambda(x).F(b)\}.$$

Now K acts via the left regular representation. It turns out that this is the finite-dimensional irreducible repn of K with highest weight λ.

5 Holomorphic discrete series.

This section will be superceded by a section on discrete series!

Let G be linear connected reductive, and write $\mathfrak{g} = \mathfrak{k} \oplus \mathfrak{p}$. Let $\mathfrak{c}$ be the centre of $\mathfrak{k}$. Now **assume that the centralizer of $\mathfrak{c}$ in $\mathfrak{g}$ is $\mathfrak{k}$.**. It certainly contains $\mathfrak{k}$. This is a non-trivial condition. It succeeds if G is compact but it implies that G has compact centre. It works for $\mathrm{sl}_2(\mathbf{R})$ but not for $\mathrm{sl}_n(\mathbf{R})$ for any $n \geq 3$. It implies that any Cartan subalgebra of $\mathfrak{k}$ is also a Cartan subalgebra of $\mathfrak{g}$ so **our assumption implies that the ranks of G and K coincide** and in particular our assumption implies that G has a compact Cartan. Our assumption holds for $G = \mathrm{SU}(m, n)$ and $\mathrm{Sp}_{2n}(\mathbf{R})$, and for compact groups, but not for $\mathrm{SL}_n(\mathbf{R})$.

Choose a Cartan $\mathfrak{h}$ in $\mathfrak{k}$, tensor up to $\mathbf{C}$, so $\mathfrak{h} \otimes \mathbf{C}$ is a Cartan in $\mathfrak{g} \otimes \mathbf{C}$, and perform the usual decomposition into 1-dimensional root spaces. Each root space is either in $\mathfrak{k} \otimes \mathbf{C}$ or $\mathfrak{p} \otimes \mathbf{C}$, because $\mathfrak{h}$ is in $\mathfrak{k}$, and our assumption means that each root is either *compact* (in $\mathfrak{k} \otimes \mathbf{C}$) or *non-compact.*

In the special case where G is compact, all the roots are also compact. If $G = \mathrm{SU}(m, n)$ with respect to the diagonal form then the compact roots are the ones that live in either the top $m \times m$ block or the bottom $n \times n$ one, so there are $m(m-1)/2 + n(n-1)/2$ compact roots and mn noncompact ones.

Choose a minimal parabolic B of $G_{\mathbf{C}}(\mathbf{C})$ using $\mathfrak{h} \otimes \mathbf{C}$, using $\theta\mathfrak{n}$ as in the Borel-Weil theorem, and for λ a weight consider consider holomorphic functions from the subgroup GB of $G_{\mathbf{C}}(\mathbf{C})$ which transform on the right by B via λ, and for which a certain integral converges—see p158 of Knapp (section VI.4). Theorem 6.6 of Knapp says that this representation is non-zero and irreducible *if* λ is dominant with respect to the compact roots and, essentially, "non-dominant" with respect to the non-compact roots! These are the holomorphic discrete series. Moreover the finite-dimensional rep of K with highest weight λ shows up in the restriction of this rep to K.

Theorem. The holomorphic discrete series are unitary and their matrix coefficients are square-integrable. [this is in Theorem 6.6 of Knapp.]

6 Principal series.

Let G be linear connected reductive. Let $S = MAN$ be a parabolic. Choose an irreducible unitary representation σ of M on a space V. Let ν be a character of A, so it's an element of $\mathrm{Hom}_{\mathbf{R}}(\mathfrak{a}, \mathbf{C})$ with $\mathfrak{a}$ the Lie algebra of A. We now induce up to get a repn $U(S, \sigma, \nu)$ of G: Consider an appropriate completion of the continuous functions $F : G \to V$ such that $F(xman) = e^{-\nu+\rho} \log(a)\sigma(m)^{-1}F(x)$, where ρ is half the sum of the roots of $(\mathfrak{g}, \mathfrak{a})$ (remember that this is not an abstract root system) which are positive with respect to the choice of N.

This space might not be irreducible but it is a Hilbert space with an action of G (G acts thus: $(gF)(x) = F(g^{-1}x)$). All the results below are in chapter VII of Knapp.

Theorem. If ν is pure imaginary (i.e. the induced rep of A is unitary) then $U(S, \sigma, \nu)$ if unitary. [remark: sometimes the induced rep can be unitarizable even if ν isn't pure imaginary.]

Theorem. If S is the upper triangular subgroup of $G := \mathrm{SL}_n(\mathbf{C})$ then the induced rep is always irreducible.

Theorem. For S minimal, the space of continuous endomorphisms of $U(S, \sigma, \nu)$ is finite-dimensional.

Theorem $U(S, \sigma, \nu)$ is admissible.

A *principal series* rep is $U(S_p, \sigma, \nu)$ with ν imaginary and S_p minimal. A *non-unitary principal series* is $U(S_p, \sigma, \nu)$ with S_p minimal but ν arbitrary. Note that S_p being minimal implies that $M = M_p$ is compact, as $\mathfrak{m} = \mathfrak{g}_0 \cap \mathfrak{k} \subseteq \mathfrak{k}$.

Theorem If S is any parabolic, if σ is irreducible tempered (see below) unitary, and if ν is such that its real part is in the open positive Weyl chamber, then $U(S, \sigma, \nu)$ has a unique irreducible quotient $J(S, \sigma, \nu)$ (the Langlands quotient). Note that this is not quite right: we should really work on the level of $(\mathfrak{g}, K)$-modules. See Theorem 7.24' in Knapp, on p214, just before Corollary 8.13.

7 Admissible reps.

I know what these are: each rep of K only shows up finitely often.

Theorem Irreducible unitary implies admissible.

Theorem $S = MAN$ and σ is irreducible unitary then $U(S, \sigma, \nu)$ is admissible.

Theorem Infinitesimally equivalent admissible reps have the same K-finite matrix coefficients.

Theorem Casselman's subrep theorem: any irred admiss rep is infinitesimally equiv with a sub rep of a nonunitary principal series.

8 Tempered reps.

If V is an irreducible admissible rep of G then V is *tempered* if all the K-finite matrix coefficients of V are in $L^{2+\epsilon}$ for all $\epsilon > 0$.

Prop 7.14 of Knapp: if S is parabolic and σ is irreducible *tempered* unitary on M and if ν is imaginary then $U(S, \sigma, \nu)$ is also tempered (and unitary).

Theorem (Knapp 8.53): For V irred admiss, V is tempered iff V is inf equiv with a subrep of a $U(S, \sigma, \nu)$ with σ discrete series and ν imaginary.

Corollary (Knapp p260) Discrete series implies inf equiv with unitary, and hence tempered implies inf equiv with unitary.

Theorem (8.54) Fix a min parabolic $S_p = M_p A_p N_p$. Then inf equiv classes of irred admiss reps of G biject with triples $(S, [\omega], \nu)$ with $S = MAN$ a parabolic containing S_p, ω irreducible unitary tempered, and ν having real part in the positive Weyl chamber.

9 Square integrable—or discrete series.

We say V is *square-integrable* if it is unitary and some non-zero matrix coefficient is in L^2. In fact for irreducible unitary representations, square-integrable is the same as all matrix coefficients being in L^2, and is also the same as V being a sub (or a direct summand) of $L^2(G)$. So square-integrable means the same thing as discrete series, it seems to me.

From another place in Knapp: if V is irred admissible then we say that it's discrete series if all its K-finite matrix coefficients are in $L^2(G)$. This is equivalent to being infinitesimally equiv with an irred subrep of the right regular rep on $L^2(G)$.