A tour through the proof of Margulis Superrigidity

Sayantan Khan
University of Michigan

Thursday 3rd October, 2019

Abstract

Margulis' Superrigidity theorem tells us that for a higher rank semisimple Lie group G, any representation of a lattice Γ extends to a continuous representation of G (under reasonably mild conditions). In this talk, we'll go through a proof of this fact, and along the way see how this proof combines ideas from Lie theory, dynamics, and random walks on groups. In particular, we'll see how a Γ-stationary measure on G/P is constructed, and combine that with measure proximality to prove Margulis superrigidity.

The Margulis Superrigidity theorem essentially tells us that representations of irreducible lattices in semisimple Lie groups extend to continuous representations of the entire Lie group. Of course, the version as stated isn't quite correct: one overkill way to see it is to consider a surface group, which is a lattice in $\mathrm{PSL}(2,\mathbb{R})$, a rank 1 Lie group. Each outer automorphism of the lattice induces a representation of the lattice into $\mathrm{PSL}(2,\mathbb{R})$. Not all of these representations can have continuous extensions, since the outer automorphism group of a surface group is infinite, whereas semisimple Lie groups have finite outer automorphism groups. A version of Margulis Superrigidity that is true is the following.

Theorem 1 (Margulis Superrigidity)**.** *Suppose we have the following.*

- G *is a semisimple Lie group (inside some* $\mathrm{SL}(n,\mathbb{R})$*) of real rank at least* 2*.*
- Γ *is an irreducible lattice in* G*.*
- H *is a connected, non-compact, simple subgroup of some* $\mathrm{SL}(m,\mathbb{R})$ *with trivial center.*
- $\varphi : \Gamma \to H$ *is an irreducible representation.*
- $\varphi(\Gamma)$ *is Zariski dense in* H*.*

Then φ *extends to a continuous homomorphism* $\widehat{\varphi} : G \to H$*.*

This is not the most general version of the theorem, but the proof of this version has all the key ingredients, and provides some insight into what's going on in the proof. The proof we'll be looking at is taken from [Mor01], and a couple of lemmas about stationary measures are from [Mar91].

The proof consists of two reductions. The first reduction involves reducing the problem of finding a continuous representation of G to finding a G-invariant finite dimensional space of continuous sections of a certain vector bundle. The next reduction involves reducing the problem to finding just a single A-invariant *measurable* section (where A is the maximal abelian subgroup).

G-invariant sections of a flat vector bundle

Let V be the vector space Γ acts on. Combining that action with right multiplication on G, we get a right action of Γ on $G \times V$, given by the following expression.

$$(x, v) \cdot \gamma := (x\gamma, \varphi(\gamma^{-1})v)$$

If we quotient out $G \times V$ by this action of Γ, we get a vector bundle $\mathcal{E}_\varphi$ over G/Γ (a vector bundle defined in this manner is said to be flat). It's easy to see that sections of this vector bundle correspond exactly to right Γ-equivariant maps from G to V, i.e. maps ξ that satisfy the following property for all $\gamma \in \Gamma$.

$$\xi(g\gamma) = \varphi(\gamma^{-1})\xi(g)$$

We can define a right action of G on sections by letting $g \cdot \xi(h) = \xi(hg)$. We would like to focus on finite-dimensional G-invariant subspaces of sections.

To see how all this helps with constructing a continuous extension $\widehat{\varphi}$, consider the following special case.

Lemma 2. *Suppose W is a finite dimensional G-invariant space of sections of $\mathcal{E}_\varphi$ such that the evaluation map sending ξ to $\xi(e)$ is a bijection. Then there exists a continuous extension $\widehat{\varphi}$ of φ.*

Proof. Note that the group G acts on W, which is a space of sections. We can turn this into an action $\widehat{\varphi}$ of G on V by merely evaluating the sections at e. One can then check that if we restrict this action to Γ, the action is exactly the same as φ, which means it's an extension of φ. $\square$

The next thing we need to do is drop the assumption that the finite dimensional G-invariant subspace W is in bijection with V. This is where we use the fact that φ is an irreducible representation. Consider W as a representation of G. Since G is semisimple, W breaks up as a sum of irreducible representations of G. Because Γ is Zariski dense in G, all of these are also irreducible Γ-representations. The evaluation map at e is an intertwining operator, which means by Schur's lemma, we must have that the evaluation map is either an isomorphism, in which case we're done by Lemma 2, or the zero map. In that case, by G-invariance, we must have that W is 0-dimensional, which violates our assumptions.

We have thus shown that to find a continuous extension $\widehat{\varphi}$, it suffices to find a non-trivial finite dimensional G-invariant space of continuous sections of $\mathcal{E}_\varphi$. Similarly, to find a measurable extension, it suffices to find a G-invariant space of measurable sections, which is what we'll actually do. Note that this doesn't affect our result, since measurable group homomorphisms are automatically continuous.

Reducing to the case of finding an A-invariant section

So far, we've managed to reduce the problem of finding a continuous/measurable extension $\widehat{\varphi}$ to just finding a finite dimensional G-invariant space of sections of $\mathcal{E}_\varphi$. Just finding a space of G-invariant sections isn't too hard: one can take any section, and look at its G-orbit. However, that doesn't guarantee that the space of sections one ends up getting is finite dimensional. To prove finite-dimensionality, we need to use some ideas from dynamics. We have the following lemma.

Lemma 3. *Suppose H is a closed, non-compact subgroup of the maximal abelian subgroup A. Let W be an H-invariant subspace of $\mathrm{Sect}(\mathcal{E}_\varphi)$ that is finite dimensional, then $\langle \mathcal{C}_G(H) \cdot W \rangle$ is finite dimensional.*

Proof. For simplicity, assume that W is 1-dimensional, and a section σ is left invariant by H. Because H is non-compact, by Moore Ergodicity, we have that the orbit of a lot of points are dense. Since the section is continuous, it is determined by its value at one point, and hence, the set of all H invariant sections is finite dimensional. Since $\langle \mathcal{C}_G(H) \cdot W \rangle$ is a subset of the H-invariant sections, this must be finite dimensional as well. $\square$

This lemma gives us a way arguing that the space of sections is finite dimensional we get when we just look at the span of the orbit under some group elements. To make it work in our case, where we want to look at the G orbit of some section, we need the following Lie theoretic lemma, which we won't prove.

Lemma 4. *If G is a semisimple Lie group of rank greater than or equal to 2, then there exist closed subgroups $\{L_1, \ldots, L_r\}$ such that the following conditions hold.*

1) $G = L_r L_{r-1} \cdots L_1$.

2) The groups $H_i := L_i \cap A$ and $H_i^\perp := \mathcal{C}_A(L_i)$ are non-compact.

Using Lemmas 3 and 4, we can find a finite dimensional space of G-invariant sections. To do so, we start with a continuous/measurable A-invariant section σ. Let V_0 be the span of σ. We iteratively define V_i to be $\langle L_i \cdot A \cdot V_{i-1} \rangle$. Clearly, V_r is the G-span of σ. To show that all the V_is are finite dimensional, we induct. Suppose V_{i-1} is finite dimensional. By the induction hypothesis, V_{i-1} is L_{i-1}-invariant, and hence H_{i-1}-invariant. Since A commutes with everything in H_{i-1}, Lemma 3 tells us that $\langle A \cdot V_{i-1} \rangle$ is finite dimensional. This finite dimensional space of sections is $H_i^\perp$-invariant, and hence L_i-invariant.

Finding an A-invariant section

Let us restate our goals now that we have made the two reductions we needed to.

Theorem 5. *Suppose we have the following.*

- *G is a semisimple Lie group (inside some $\mathrm{SL}(n, \mathbb{R})$) of real rank at least 2.*
- *Γ is an irreducible lattice in G.*
- *H is a connected, non-compact, simple subgroup of some $\mathrm{SL}(m, \mathbb{R})$ with trivial center.*
- *$\varphi : \Gamma \to H$ is an irreducible representation.*
- *$\varphi(\Gamma)$ is Zariski dense in H.*

Then there exists a non-zero A-invariant measurable section of $\mathcal{E}_\varphi$, or equivalently, a right Γ-equivariant measurable function from G/A to $\mathbb{R}^n$, which is the vector space on which H acts.

To prove the reduced version of Margulis Superrigidity, we'll do something quite strange. Rather than immediately focusing on G/A, we'll focus on G/P, where P is the minimal parabolic subgroup. When G is just $\mathrm{SL}(n, \mathbb{R})$, P is the space of upper triangular matrices. The reason why we look at G/P is the following: this space can be identified with the space of flags in $\mathbb{R}^n$, for $\mathrm{SL}(n, \mathbb{R})$. The space $G/P \times G/P$ can be identified with pairs of flags in $\mathbb{R}^n$. We can now look at the subspace of pairs of flags that lie in general position: this subspace forms a full measure subset of $G/P \times G/P$, and G acts transitively on this space. Furthermore, the stabilizer of such a pair of flags in general position is exactly A, which means that up to a set of measure 0, $G/A \cong G/P \times G/P$. This explains why we need to find a measurable section rather than a continuous one.

We now just need to find a right Γ-equivariant map from $G/P \times G/P$ to $W = \mathbb{R}^n$. It might help to focus on one factor at a time. Let's try to find a Γ-equivariant map from G/P to W. It might help to note that P is a solvable subgroup of G, hence amenable. That means we can try and use Furstenburg's lemma, which states the following.

Lemma 6. *If P is a closed, amenable subgroup of G, and Γ acts continuously on a compact metric space X, then there exists a Borel measurable map from G/P to $\mathrm{Prob}(X)$ which is Γ-equivariant, except on a set of measure 0.*

Let's see how we can use this lemma in our situation. We have an action of Γ on W, but that certainly isn't a compact metric space. However, $\mathbb{P}(W)$ is a compact metric space on which Γ acts. This means we get a map $\widehat{\xi} : G/P \to \mathrm{Prob}(\mathbb{P}(W))$ which is Γ-equivariant. This clearly isn't good enough, since we wanted a map into W, and not $\mathrm{Prob}(\mathbb{P})(W)$. However, if we could upgrade this map into a map ξ from G/P to $\mathbb{P}(W)$, that would suffice, since we could consider the dual map from G/P to $\mathbb{P}(W^*)$, and take the tensor product of those maps. That would give us a map into the following space.

$$\xi \otimes \xi^* : G/P \times G/P \to \mathbb{P}(W \otimes W^*) \cong \mathbb{P}(\mathrm{End}(W))$$

We could lift this map up to a map to just $\mathrm{End}(W)$ by fixing the trace to be 1 (there is some work involved in showing that this only leaves out a set of measure 0). Once we have that, we get a representation of G acting on $\mathrm{End}(W)$, but we can then look at each of the irreducible components separately, and get what we want.

That means most of the hard work of the proof is really concentrated in upgrading the map $\widehat{\xi} : G/P \to \mathrm{Prob}(\mathbb{P}(W))$ to a map $\xi : G/P \to \mathbb{P}(W)$. One reasonable way of upgrading $\widehat{\xi}$ would be to show that its image lies in the set of Dirac measures up to a set of measure 0.

We do this in the most absurd way possible: we pick a random sequence of elements γ_i in the group Γ. We apply these elements to a point p in G/P. Since G/P is compact, after passing to a subsequence, we get a convergent sequence. Also, since the action is equivariant, we can look at the orbit of the corresponding point in $\mathrm{Prob}(\mathbb{P}(W))$, and if we can show that such a sequence converges to a Dirac mass, then we'll be done.

In short, we'll look at the orbit of a probability measure in $\mathrm{Prob}(\mathbb{P}(W))$ and show that it converges almost surely to a Dirac mass. To do this, we'll prove that the action of Γ on $\mathbb{P}(W)$ is *proximal.*

Lemma 7 (Proximality of the Γ-action)**.** *For any $[w_1]$ and $[w_2]$ in $\mathbb{P}(W)$, there exists a sequence $\{\gamma_n\}$ in Γ such that $d(\gamma_n[w_1], \gamma_n[w_2])$ converges to 0.*

Proof. Pick a diagonalizable element $\gamma \in \Gamma$ which has a unique eigenvalue of maximal absolute value. Without loss of generality, we can assume it's a simple eigenvalue, otherwise, we just look at the appropriate exterior power. In the case of $\mathrm{SL}(2, \mathbb{R})$, this corresponds to a hyperbolic element in Γ. We form a basis using the eigenvalues of γ, and denote the largest eigenvector by v_1. Clearly, for any two vectors w_1 and w_2 with a non-zero v_1 components, $\gamma^n[w_i]$ converges to $[v_1]$. For something with a zero v_1 component, we multiply both w_1 and w_2 with some δ such that δw_i has a non-zero v_1 component, and then the result follows. □

We can in fact strengthen the previous argument for measures on $\mathbb{P}(W)$.

Lemma 8 (Measure proximality)**.** *Let μ be any measure on $\mathbb{P}(W)$. Then there exists a sequence $\{\gamma_i\}$ in Γ such that $\gamma_i^* \mu$ converges to a Dirac mass.*

In particular, the above lemma also shows that there is no Γ-invariant probability measure on $\mathbb{P}(W)$. However, if we pick a probability measure ν on Γ, then there is a probability measure μ on $\mathbb{P}(W)$ which is invariant on average. To be more precise, it satisfies the following identity.

$$\sum_{\gamma \in \Gamma} \nu(\gamma)\gamma^* \mu = \mu$$

What we do now is perform a random walk on Γ with respect to the probability ν, and look at the orbit of μ under this random walk. As it turns out, for almost every random walk, the orbit of μ goes to some Dirac mass.

Theorem 9 (Mean proximality)**.** *Assume that ν is supported on all of Γ, and μ is a ν-stationary probability measure on $\mathbb{P}(W)$. Then for almost every random walk, $(\gamma_1 \cdots \gamma_n)^* \mu$ converges to some δ_c.*

Proof. First of all, note that the sequence of random variables $(\gamma_1 \cdots \gamma_n)^* \mu$ forms a bounded martingale, which means it converges almost surely. Thus, to show that it converges to a Dirac mass, it suffices to show convergence along a subsequence. To do this, pick the sequence $\{g_i\}$ such that $g_i^* \mu$ converges to a Dirac mass. We would like to say that almost every random walk contains a subsequence whose last element is g_i. This happens almost surely because all the g_i have positive ν-mass. If we look at the orbit of μ under this subsequence, it converges to some point mass, which proves our claim. □

Now we have everything we need to prove our claim rigourously. Consider a Lebesgue class measure μ on G/P which is ν-stationary. We look at the following measure on $\mathbb{P}(W)$.

$$\mu_{\mathbb{P}(W)} = \int_{G/P} \widehat{\xi}(x) d\mu(x)$$

This is clearly ν-stationary, since μ is. We look at the distance between $(\gamma_1 \cdots \gamma_n)^* \mu_{\mathbb{P}(W)}$ and $\delta_{\mathbb{P}(W)}$. By means proximality, this distance goes to 0. By the virtue of it being ν-stationary, we see that this distance equals the distance between $\widehat{\xi}(x)$ and $\delta_{\mathbb{P}(W)}$, which means it must be 0. This proves our result, and by reduction, concludes the proof of Margulis Superrigidity.

References

[Mar91] Gregori A Margulis. *Discrete subgroups of semisimple Lie groups*, volume 17. Springer Science & Business Media, 1991.

[Mor01] Dave Witte Morris. Introduction to arithmetic groups, 2001.