Geometric group theory
Lecture 1

Lawk Mineh

October 2025

1 The basics

1.1 Group presentations

We will begin with the notion of a group presentation, which is a fundamental way to express an abstract (discrete) group. A group presentation is a description of a group in terms of 'generators' and 'relations'. That is, some free variables and the equations bounding them. To make this notion precise, we recall the free group.

Definition 1.1 (Free group)**.** Let S be a set. The *free group generated by* S is the group $F(S)$ such that for any group G and function $f\colon S \to G$, there is a unique homomorphism $\widehat{f}\colon F(S) \to G$ making the following diagram commute

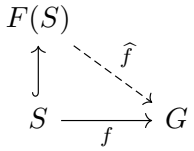

where $S \to F(S)$ is a natural inclusion.

In other words, a free group on a set S is the image of S in **Grp** under the free functor **Set** $\to$ **Grp**.

Exercise 1.2. Show that $F(S)$ and $F(T)$ are isomorphic if and only if S and T are in bijection. It follows that a free group is uniquely determined up to isomorphism by the cardinality of its generating set.

Definition 1.3 (Rank of a free group)**.** Let S be a set. The *rank* of $F(S)$ is the cardinality $|S|$ of S.

It is often useful to have a practical model of the free group that one can refer to, when it is unwieldy or otherwise not possible to use the definition above in terms of a universal property.

Definition 1.4. Let S be a set. Denote by S^{-1} the set in bijection with S, whose elements are the symbols s^{-1} for each $s \in S$: these are the formal inverses of elements of S. We identify this bijection $\cdot^{-1}\colon S \to S^{-1}$ with its inverse, so that we may write $(s^{-1})^{-1} = s$. A *word in* S is an ordered finite sequence of elements in $S \cup S^{-1}$; the *empty word* is the empty sequence. The *length* $\ell(w)$ of a word w is the number of terms in the sequence. A word is called *reduced* if it contains no consecutive terms of the form ss^{-1} for $s \in S \cup S^{-1}$. Define the equivalence relation $=_{F(S)}$ on words as the symmetric and transitive closure of deleting such an element-inverse pair. Note that any word is equivalent to a reduced word.

The free group $F(S)$ on S is the set of words in S up to the above equivalence relation, with the operation of concatenation of (class representatives) of words. It is straightforward to check that this is a well-defined operation. The identity of this group is the empty word.

It will usually not cause confusion for us to identify words in S with their equivalence classes up to reduction, so we will interchangeable refer to words as 'being' elements of $F(S)$ as well as 'representing' elements of $F(S)$. Note that every group is the quotient of a free group: indeed, if G is a group, then applying the forgetful functor it can be viewed as a set, and the universal property implies there is a unique homomorphism $F(G) \to G$ acting as 'evaluation' of words in G.

Definition 1.5 (Group presentation)**.** Let S be a set, and R a set of words in S. Let G be the quotient group $F(S)/\langle\langle R \rangle\rangle$ and write

$$\langle S \,|\, R \rangle$$

for the *presentation of* G with *generators* S and *relators* R. If both S and R are finite sets, then the presentation is called *finite*. If a word w in S represents an element $g \in G$, we may write $w =_G g$.

We say that G is *finitely generated* if it admits a presentation $\langle S \,|\, R \rangle$ with S finite. Equivalently, it is the quotient of a finite rank free group. Further, G is *finitely presented* if it admits a finite presentation.

Exercise 1.6. Show that the group of invertible $n \times n$ integer matrices $\mathrm{GL}_n(\mathbb{Z})$ is finitely generated. Show that the rational numbers $\mathbb{Q}$ do not form a finitely generated group under addition.

Exercise 1.7. Show that the properties of being finitely generated and being finitely presented are stable under extensions.

When we have explicit sets to work with, we often write a group presentation with the elements of S and R, omitting the set brackets.

Example 1.8.

- The free group on S has a presentation with generating set S and no relators;
- the free abelian group on S has a presentation with generators S and commutators $[s, t]$ as relators for each $s, t \in S$;

- the fundamental group of the genus g surface Σ has presentation

$$\langle a_1, \ldots, a_g, b_1, \ldots, b_g \,|\, [a_1, b_1] \ldots [a_g, b_g] \rangle.$$

  To see this, observe that Σ can be obtained as the quotient space of a $4g$-gon. One should convince oneself first about the case that Σ the torus $g = 1$, then observe that higher genus surfaces are obtained by taking connected sums of tori and lower genus surfaces. The $4g$-gon can be cut into g hexagons, each of which are tori with a single boundary component under the edge identifications;
- a cyclic group of order n has presentation $\langle a \,|\, a^n \rangle$.

Of course, a presentation does not determine a group uniquely. Indeed, though some information can be gleaned from a presentation in specific circumstances, group presentations in general do not encode readily accessible information about a group. This can be seen more concretely from the fact that there is no algorithm that, given a group presentation, can determine even whether the group it describes is the trivial group. However, if we know which group we are working with to start off with, two of its presentations are not *too* unrelated (at least, when it comes to finite presentations).

Definition 1.9 (Tietze transformations)**.** Let $\mathcal{P} = \langle S \,|\, R \rangle$ be a presentation for a group G. The following four operations are *Tietze transformations*, taking the presentation $\mathcal{P}$ to a presentation $\mathcal{P}'$:

(i) Let $r \in \langle\langle R \rangle\rangle$ be a word in $F(S)$. Define $\mathcal{P}' = \langle S \,|\, R \cup \{r\} \rangle$.
(ii) Suppose $r \in R$ is such that $r \in \langle\langle R - \{r\} \rangle\rangle$. Define $\mathcal{P}' = \langle S \,|\, R - \{r\} \rangle$.
(iii) Let t be an element in $F(S)$ and w a word in S representing t. Define $\mathcal{P}' = \langle S \cup \{t\} \,|\, R \cup \{t^{-1}w\} \rangle$.
(iv) Suppose $s \in S$ is such that s can be written as a word w in $S' \subset S$, and $s^{-1}w \in R$. Define $\mathcal{P}' = \langle S - \{s\} \,|\, R - \{s^{-1}w\} \rangle$.

These operations correspond to adding a superfluous relator, deleting a superfluous relator, adding a superfluous generator, and deleting a superfluous generator respectively.

It is a tedious, though possibly instructive, exercise to verify that each of the Tietze transformations preserve the isomorphism type of the presented group G. We now have the important observation about finite presentations.

Lemma 1.10 (Tietze's theorem)**.** *Let $\mathcal{P}$ and $\mathcal{P}'$ be two finite presentations of a given group G. Then there exist a finite sequence of Tietze transformations that transform $\mathcal{P}$ into $\mathcal{P}'$.*

Proof. The idea is that one can arrive at a common presentation for both $\mathcal{P}$ and $\mathcal{P}'$ by adding in all of the generators and then relators from both presentations in one at a time. We leave the details to the reader. □

1.2 Groups and their actions

Recall that an *action* of a group G on a set X is a map $\cdot \colon G \times X \to X$ such that

- $1 \cdot x = x$ for all $x \in X$; and
- $g \cdot (h \cdot x) = (gh) \cdot x$ for all $g, h \in G$ and $x \in X$.

Let us introduce some terminology for group actions.

Definition 1.11 (Group actions)**.** Let X be a set with a G-action $\cdot\colon G \times X \to X$. We say the action is:

- *free* if $g \cdot x = x$ implies $g = 1$;
- *transitive* if for any $x, y \in X$, there is $g \in G$ with $g \cdot x = y$;
- *faithful* if every element if for every $g \in G$, there is $x \in X$ with $g \cdot x \neq x$;

when X is a topological space:

- *cocompact* if the quotient space X/G is compact, with the quotient topology;
- *properly discontinuous* if for any compact $K \subseteq X$, the set

$$\{g \in G \,|\, gK \cap K \neq \emptyset\}$$

  is finite;

and when X is a metric space, with metric d:

- *by isometries* if for any $g \in G$ and $x, y \in X$, we have

$$\mathrm{d}(g \cdot x, g \cdot y) = \mathrm{d}(x, y)$$

- *cobounded* if the quotient space X/G is bounded;
- *geometric* if it is by isometries, cocompact, and properly discontinuous.

It is often useful to think about an action of a group G on a set X as a homomorphism $G \to \mathrm{Aut}(X)$ from the group to the automorphisms of the set. We say automorphisms here, rather than permutations, for X is often endowed with some structure, and the action is required to preserve that structure. For instance, if X is a topological space, we want the group elements to act by homeomorphisms, when X is a metric space, usually by isometries, and so on.

We now introduce the most basic geometric-combinatorial object for our study: *Cayley graphs* of groups. For us, a *graph* Γ will be a set of *vertices* $V\Gamma$ and a set of *edges* $E\Gamma$, which comes equipped with a pair of functions $\iota, \tau\colon E\Gamma \to V\Gamma$ (denoting the *initial* and *terminal* endpoints of an edge). We will say that two vertices $v, w \in V\Gamma$ are *connected* by an edge $e \in E\Gamma$ with $\iota(e) = v$ and $\tau(e) = w$, and we write $v \sim w$ in this case.

The *geometric realisation of a graph* Γ is a simplicial complex whose 0-skeleton is $V\Gamma$, and whose 1-simplices are the edges $E\Gamma$, with attaching maps determined by the incidence functions ι and τ. We equip this complex with the metric induced by giving each edge unit length. Throughout, we will identify a graph with its geometric realisation.

Definition 1.12 (Cayley graph)**.** Let G be a group with generating set S. The *Cayley graph* of G with respect to S is the graph $\Gamma(G, S)$ whose vertex set is G, and with an edge $g \sim h$ if there is $s \in S$ with $gs = h$.

Note that G acts (by left multiplication) transitively on the vertex set of $\Gamma(G, S)$, and with $|S|$-many orbits of edges. It is straightforward to see that this is an isometric

action. When the group G is finitely generated and S is a finite set, there are finitely many edge orbits, and so the action is also properly discontinuous and cocompact. This gives us our archetypal model for a geometric action. We will see later that essentially every geometric action of a finitely generated group is like one on a Cayley graph.

Definition 1.13 (Word metric)**.** Let G be a group and S be a generating set. The *word metric* on G with respect to S is the metric d_S defined at

$$\mathrm{d}_S(g,h) = \min\{\ell(w) \,|\, w =_G g^{-1}h\}.$$

We will write $|g|_S = \mathrm{d}_S(1,g)$ for the *length* of g with respect to S.

The word metric on a group coincides with the restriction of the edge-path metric on the associated Cayley graph to its vertex set, and the length of an element with respect to a generating set is exactly the length of the shortest word representing that element in that generating set.

Example 1.14. The free abelian group $\mathbb{Z}^n = \langle a_1, \ldots, a_n \,|\, [a_i, a_j] = 1\rangle$ acts geometrically on the Euclidean space $\mathbb{R}^n$ of dimension n, by the translations

$$a_i \cdot (x_1, \ldots, x_n) = (x_1, \ldots, x_i + 1, \ldots, x_n).$$

This is, by construction, an action by isometries, and one should check that the action is properly discontinuous. The quotient space $\mathbb{R}^n/\mathbb{Z}^n = \mathbb{T}^n$ is the Euclidean n-torus, the product of n copies of the circle S^1. One can view this as the covering space action of $\pi_1(\mathbb{T}^n) \cong \mathbb{Z}^n$ on its universal cover $\mathbb{R}^n$.

Taking $S = \{a_1, \ldots, a_n\}$ as the standard generating set as in the presentation above. The Cayley graph $\Gamma = \Gamma(\mathbb{Z}^n, S)$ embeds as the integer grid in $\mathbb{R}^n$, and the action of $\mathbb{Z}^n$ above restricts to the standard action on Γ by left multiplication. Note that this embedding of Γ into $\mathbb{R}^n$ is not quite isometric: one can show it is distance non-increasing, and decreases distances by at most a factor of $\sqrt{n}$.

Example 1.15. Let S be a set, and $F = F(S)$ the free group on S. The Cayley graph $\Gamma(F, S)$ is the regular tree with valence $2|S|$: each vertex has an outgoing edge labelled s and an incoming edge with labelled s^{-1}. Note that F is the fundamental group of the wedge of $|S|$ circles, whose universal cover is $\Gamma(F, S)$.

Example 1.16 (Cyclic group)**.** An infinite cyclic group $\langle a\rangle \cong \mathbb{Z}$ is a free group, and has an obvious one-generator presentation with no relators. As a metric space, the Cayley graph with respect to this generating set is, of course, just the line $\mathbb{R}$. Consider, however, the generating set $S = \{a^2, a^3\}$ for $\langle a\rangle$. The Cayley graph $\Gamma = \Gamma(\mathbb{Z}, S)$ is definitely not a line: it has many loops for instance.

However, the map $\Gamma \to \mathbb{R}$ defined by taking the identity on $V\Gamma = \mathbb{Z}$ and mapping each edge to its numerically lesser endpoint distorts distances additively by at most 3. In this way, Γ is 'coarsely' isometric to the real line.

Exercise 1.17. Draw the Cayley graph for the group with presentation

$$\langle a, t \,|\, tat^{-1} = a^2 \rangle.$$

This is an example of a *Baumslag-Solitar group*: these form a two-parameter family of groups, indexed by integers $m, n \in \mathbb{Z}$, with a relator $ta^m t^{-1} = a^n$.

1.3 Quasi-isometries and the Milnor-Schwarz lemma

The examples of the previous section serve to illustrate a key point: though the exact metric on two spaces with a geometric G-action may differ on a local level, the large-scale 'rough' geometry of the spaces remains the same. The takeaway is that, for our purposes, isometry is not the correct notion of morphism, motivating the following.

Definition 1.18 (Quasi-isometry)**.** Let X and Y be metric spaces, $\lambda \geq 1$, and $c \geq 0$. A map $f\colon X \to Y$ is a (λ, c)-*quasi-isometric embedding*

$$\frac{1}{\lambda} \mathrm{d}_X(x, x') - c \leq \mathrm{d}_Y(f(x), f(x')) \leq \lambda \, \mathrm{d}_X(x, x') + c$$

for all $x, x' \in X$.

Moreover, f is K-*coarsely surjective* if for every $y \in Y$, there is $x \in X$ with $\mathrm{d}_Y(f(x), y)$. A (λ, c)-*quasi-isometry* is a (λ, c)-quasi-isometric embedding that is K-coarsely surjective for some $K \geq 0$.

As with many definitions we will give, we will omit the constants when they are not important to the discussion. That is, for example, we say a map f is simply a *quasi-isometry* if there are some $\lambda \geq 1$ and $c \geq 0$ such that it is a (λ, c)-quasi-isometry. It is important to note that a quasi-isometry need not be continuous!

Exercise 1.19. Show that if $f\colon X \to Y$ is a quasi-isometry, it has a *quasi-inverse*: a quasi-isometry $g\colon Y \to X$ such that $g \circ f$ and $f \circ g$ are a finite distance from the identity on X and Y respectively (with respect to the supremal metric on functions), in such a way that the all the constants involved depend only on those of f (i.e. they are independent of the particular function).

Show that quasi-isometry is an equivalence relation on metric spaces.

Example 1.20. Given $n \geq 2$, write T_n for the n-regular tree. For any $m, n \geq 3$, the trees T_m and T_n are quasi-isometric. Transitivity of quasi-isometry means that it suffices to show that any n-regular tree is quasi-isometric to T_3. The key fact is that if one collapses an edge of T_3, then one combines two vertices and increases the valence by one.

Thus, let us take a spanning forest $\mathfrak{T}$ of T_3 by disjoint paths of length $n-3$ (note: one needs the axiom of countable choice for this), and consider the map $f\colon T_3 \to T_3/\mathfrak{T} \cong T_n$ obtained by collapsing each connected component of $\mathfrak{T}$ to a single point. The map f is of course distance non-increasing. Moreover, at least every $(n-2)^{\text{th}}$ edge of an arc in T_3 must lie outside of $\mathfrak{T}$, so that

$$\mathrm{d}_{T_3}(x, y) \leq (n-2) \, \mathrm{d}_{T_n}(f(x), f(y)) + (n-3)$$

for any $x, y \in T_3$. Finally, our map is surjective, so it is a quasi-isometry.

Example 1.21. We will give a sketch that the real line $\mathbb{R}$ and the ray $[0,\infty)$ are not quasi-isometric, and leave it to the reader to assemble the details. Indeed, suppose that $\varphi\colon \mathbb{R} \to [0,\infty)$ is a quasi-isometric embedding, and suppose for simplicity that $\varphi(0) = 0$. As $t \to \infty$, we must have $\varphi(t) \to \infty$ and $\varphi(-t) \to \infty$, since φ coarsely preserves distances. Pick some very large $x \in [0,\infty)$. There must then be some correspondingly very large $s, t > 0$ so that $\varphi(-s)$ and $\varphi(t)$ lie a uniformly bounded distance from x. Using the fact that φ is a quasi-isometric embedding, $t - (-s) = t + s$ is uniformly bounded. Provided one picks x large enough, then $t + s$ may take arbitrarily large values: a contradiction.

Exercise 1.22.

1. Show that the quarter-plane $\mathbb{R}^2_{++} = \{(a,b) \in \mathbb{R}^2 \,|\, a, b > 0\}$ is quasi-isometric to the real ray $\mathbb{R}_+ = \{a \in \mathbb{R} \,|\, a > 0\}$, but that neither are quasi-isometric to the half-plane $\mathbb{R}^2_+ = \{(a,b) \in \mathbb{R}^2 \,|\, b > 0\}$.
2. Show that $\mathbb{R}^n$ is not quasi-isometric to T_m for any $m, n \geq 2$.
3. Show that if $\mathbb{R}^m$ is quasi-isometric to $\mathbb{R}^n$, then $m = n$. (Hint: Consider the volume growth rate of balls in $\mathbb{R}^n$. That is, the growth rate of the function $\mathrm{vol}_n\colon r \mapsto \mathrm{vol}(B_{\mathbb{R}^n}(0,r))$, which is polynomial in degree n. Show that this is rate is – up to a suitable equivalence relation – preserved under quasi-isometries.)

One of the immediate upshots of this notion is that the quasi-isometry type of the Cayley graph of a finitely generated group is an invariant of the group.

Lemma 1.23. *Let S and T be finite generating sets for group G. Then (G, d_S) and (G, d_T) are quasi-isometric.*

Proof. It is enough to show that (G, d_S) is quasi-isometric to $(G, \mathrm{d}_{S\cup T})$. The result then follows by transitivity of quasi-isometry. We may thus suppose without loss of generality that $S \subset T$, so $S = \{s_1, \ldots, s_n\}$ and $T = S \sqcup \{t_1, \ldots, t_m\}$, with $m \geq 1$.

The identity map is our candidate for a quasi-isometry. Of course, this map is surjective. It is immediate that the identity is distance non-increasing: a word in S representing an element of G must be at least as long as one in T. As S is a generating set, each of the elements t_i can be expressed as a word w_i in S. Let $\lambda = \max\{\ell(w_i) \,|\, i = 1, \ldots, m\} < \infty$ be the maximum over the lengths of these words.

Take any $g, h \in G$, and let w be a word of minimal length in T representing $g^{-1}h$. We may replace any instance of t_i in w with the word w_i to obtain a word w' in S representing $g^{-1}h$. Since we are, at worst, replacing each letter with λ letters, we have

$$\mathrm{d}_S(g,h) \leq \ell(w') \leq \lambda \ell(w) = \lambda\, \mathrm{d}_T(g,h).$$

It follows that the identity map is a $(\lambda, 0)$-quasi-isometry. □

In light of the above, we will say from now on that a finitely generated group G is *quasi-isometric* to a space X if, when equipped with a word metric with respect to a finite generating set, it is quasi-isometric to X. The above lemma shows that this notion is well-defined up to change of finite generating set.

Recall that two groups are *commensurate* if they contain isomorphic subgroups of finite index.

Lemma 1.24. *Finitely generated commensurate groups are quasi-isometric.*

Proof. It is enough to show that a group is quasi-isometric to any of its finite index subgroups. Let $H \leqslant_f G$ and let S be a finite generating set for G that contains a transversal of H. As G is finitely generated, so is H; let T be a finite generating set, without loss of generality a subset of S. The inclusion map $(H. \mathrm{d}_T) \to (G, \mathrm{d}_S)$ is then an isometric embedding, and is 1-coarsely surjective. □

Exercise 1.25. Think about what the Cayley graph of a group looks like with respect to a generating set that gives a finite presentation. Come up with a graph-theoretic characterisation of a group being finitely presentable, and use this to show that finite presentability is a quasi-isometry invariant.